

By Charlie Armstrong Adams

Premier calls for
Royal Commission

The Bondi warnings

THE BONDI TERROR ATTACK

Warnings, Blind Spots,
and the Case for a Royal Commission



This is the official Facebook page of Charlie Armstrong Adams. For updates and full content, visit my website and YouTube channel.

Website: <https://www.derpunkt.news/>



Support is optional



Otherwise, please enjoy the content for free—sharing the information is the main thing.

The Bondi Terror Attack: Warnings, Blind Spots, and the Case for a Royal Commission

By Charlie Armstrong Adams

Preface

This book is written in the space where **grief** collides with **governance**. A mass-casualty attack is never only an act of violence; it is also an audit of whether the State can still guarantee the basic conditions of public life—safety, order, and equal protection—in the open. Bondi matters precisely because it is not an isolated or easily securitised venue. It is an emblem of ordinary Australia in public view, a civic commons where different communities coexist without needing to negotiate terms of belonging. When that space is violated, the damage is not only physical; it is **symbolic**, and the symbolic wound quickly becomes political and institutional.

From the outset, the most important question has not been whether tragedy occurred—everyone can see that it did—but whether the official system that exists to prevent or blunt such events functioned as Australians have been led to believe. In the immediate aftermath, governments speak in the language of solidarity and resolve. Agencies speak in the language of “ongoing investigations,” “operational security,” and the limits of what can be said. These reflexes are not necessarily cynical; they are often institutional instinct under stress. Yet a familiar dynamic follows: where the public suspects that warnings existed, or that risk indicators accumulated, a refusal to disclose the decision chain does not calm the situation. It inflames it. Trust does not recover through assurance. It recovers through **demonstration**—through a process that can reconstruct what was known, what was shared, what was requested, what was decided, what was resourced, and why.

A core theme in this manuscript is that modern counterterrorism is as much about **framing** as it is about facts. After any attack, the State does not only investigate what happened; it also participates—deliberately or not—in deciding how the event will be understood. Labels applied early tend to harden into public memory, and those labels have consequences. In the Bondi context, the book tracks two narratives operating in parallel: one emphasising **antisemitic targeting**, the other emphasising **ISIS-inspired ideology**. Both can be simultaneously true, yet they pull accountability in different directions. One concentrates attention on protection obligations to a targeted community, the other on counterterrorism architecture, watchlists, disruption powers, and intelligence sharing. The choice of emphasis can change what remedies feel “rational,” which failures appear most salient, and what questions are treated as primary versus secondary.

The central claim of this book is not that any particular institution is guilty of intentional wrongdoing. The central claim is narrower, stronger, and more defensible: where contested chains of responsibility exist—community warnings, event policing decisions, intelligence triage, and inter-agency information flow—the public cannot reasonably be asked to settle for conclusions that arrive without compelled disclosure. This is why the debate about a **Royal Commission** is not a procedural sideshow. It is the mechanism by which a democracy decides whether it will accept an official narrative on trust, or whether it will require the State to earn trust through evidence.

That argument becomes even sharper once you confront what this book treats as the most corrosive driver of post-attack distrust: the “closure problem,” and the public shock that follows when prior intelligence interest becomes known only after blood has been spilled. Intelligence services operate under triage and legal thresholds, and they close lines of inquiry

for reasons that can be rational at the time. But once an attack occurs, closure looks retroactively like evasion or incompetence—regardless of whether either is true. This mismatch in risk philosophy is predictable: agencies triage for probability under constraints; citizens triage for catastrophe and expect a different posture when the potential outcome is mass casualty.

That is why **Chapter 17** is positioned as a hinge point in the book. “The Watchlist Paradox” is not simply an argument about one individual or one decision. It is an argument about how legitimacy fractures when the public learns that an alleged perpetrator had already crossed the threshold of state attention, and yet the system’s closure and re-opening logic did not prevent the later escalation. It is also an argument about standards of proof. Suspicion is easy, especially when fragments of information appear to align into a coherent story. Proof is harder, and a responsible author has to keep those categories distinct. The function of Chapter 17 is to capture the public’s commonsense inference—“surely prior interest should have mattered”—without collapsing that inference into an unproven allegation of intent. It treats closure not as an administrative endpoint but as a risk decision that must be auditable, revisitable, and accountable, because closure decisions are precisely where future disasters are either quietly prevented—or quietly incubated.

Chapter 18 then addresses the structural question that follows immediately once the watchlist paradox is in view: if the Commonwealth agencies potentially most implicated by the public’s suspicions are **ASIO** and the **AFP**, what happens to accountability when the inquiry architecture is split along jurisdictional lines? This chapter is not written to deny the value of a NSW process. A NSW Royal Commission may be powerful in examining NSW policing posture, event resourcing, and state-level ministerial decision chains. The point is about the *limits* of what a state-anchored coercive process can do when the hardest questions run into Commonwealth systems, secrecy regimes, and federal accountability pathways. A “NSW Royal Commission + federal review” arrangement can be defended publicly as efficient and cohesive. But it can also function—intentionally or not—as a **jurisdictional shield** that reduces the intensity of scrutiny applied to Commonwealth agencies compared with what would occur under a federal or genuinely joint coercive inquiry.

That is not a conspiracy claim; it is a structural observation about incentives and institutional self-protection in federations. Reviews, even independent ones, often operate behind closed doors because their evidentiary substrate is classified, and they tend to report in summaries that cannot fully disclose the decision chain. Royal Commissions, while capable of closed sessions when required, are designed to do their work publicly wherever possible and to test narratives through compelled documents and sworn evidence. Where trust is already compromised, the distinction between **assertion** and **demonstration** is everything.

This brings the reader to the political heart of the matter: why governments resist Royal Commissions even after national trauma. The book recognises the stated reasons—speed, cohesion, secrecy—without pretending those considerations are imaginary. But it also treats those reasons as incomplete. In a high-stakes failure environment, the refusal of a coercive public process can also reflect narrative management, damage control, precedent avoidance, and the desire to keep the most reputation-critical questions inside controlled formats. Whether or not that motive can be proven in any given case, it is consistent with how power behaves after crisis, and it is precisely why coercive, independent fact-finding exists as a democratic instrument.

The reader should understand the method this book tries to follow. It distinguishes clearly between (a) what is publicly reported and officially stated, (b) what can be inferred as a plausible failure mode, and (c) what would be required to substantiate stronger claims—especially claims of Machiavellian permissiveness. The book does not ask the reader to adopt a single theory. It asks the reader to see why the institutional design of inquiry matters, and why—when warning indicators, contested narratives, and cross-jurisdiction responsibility are in play—only a process with real compulsion power can plausibly restore **public trust**.

If this manuscript succeeds, it will do so not by offering certainty where certainty is unavailable, but by clarifying the question that Australia must answer after Bondi: will the State be permitted to investigate itself through controlled processes and partial disclosure, or will it be required to submit to the kind of **open, evidence-compelling** scrutiny that democracies reserve for moments when legitimacy is at stake?



Chapter 1 — When the Tail Wags the Dog: The Finlay–Albanese Dispute and the Question of Who Oversees the Overseers

A public disagreement between Australia’s **Human Rights Commissioner** and the **Prime Minister** after the **Bondi Junction attack** is not a side story. It is an unusually clear window into the real contest that follows modern terrorism: not only the contest over **facts**, but the contest over **authority**—who is permitted to **define** what happened, who may **compel** what evidence, and who gets to decide when **scrutiny** has gone far enough. When **independent statutory officers** and **elected leaders** clash over whether the nation’s highest form of inquiry should be triggered, the dispute is never only about **procedure**. It is about the **constitutional hierarchy** of the **security state**.

The protagonists matter because of what their offices represent. The **Human Rights Commissioner’s** function is to insist, especially in moments of **national trauma**, that emergency politics does not quietly become permanent governance by **assurance** and **secrecy**. The **Prime Minister’s** function is to lead the **executive**, including the **security**

apparatus, and to select an **accountability mechanism** that can credibly restore **public trust** while protecting genuinely **sensitive capabilities**. The dispute becomes politically significant when those functions diverge in public view—when one side presses for **coercive, independent scrutiny** and the other side answers, in effect, that **internal confidence** and **targeted reviews** are sufficient.

The Commissioner's call for a **Royal Commission** is best understood as a claim about the kind of **truth** the public is entitled to after **mass casualty**: not a narrative, not a reassurance, and not a curated summary, but a **reconstructable decision chain**. Bondi, on any responsible reading, raises questions that do not resolve themselves through ordinary political messaging: whether **warning signs** existed and were properly weighted; whether information moved across **agencies** and **jurisdictions** in time to matter; whether **operational choices** about event protection were defensible; and whether the legal and governance framework is fit for a threat environment where lone actors can convert ideology into **rapid mass killing**. These are not merely security questions; they are **human-rights** and **legitimacy** questions, because the right at issue is the most basic one a State claims to protect in the open—**security of the person, equal protection, and safety** in public life.

The Prime Minister's rejection of a **Royal Commission**, and his reliance on **security** and **intelligence leadership** to support that rejection, is best understood as a counter-claim about governance under pressure. It says: there are already processes in motion; duplication risks **delay**; the wrong format can turn intelligence evidence into political **spectacle**; and sensitive material does not belong in the broadest and most public forum available. In the abstract, these reasons can sound **responsible**, even unavoidable. The problem is not that the executive is wrong to care about speed or secrecy. The problem is the institutional **optics** and **incentives** created when the executive's justification rests on the preferences of agencies that might themselves fall within the **perimeter of scrutiny**.

That is where the disagreement becomes a test of **democratic hierarchy**. **Intelligence** and **security agencies** exist to serve the elected government, not to function as co-authors of the **accountability architecture** that will be applied to their own performance. They are necessarily **secretive**, necessarily **reputation-sensitive**, and necessarily trained to minimise the exposure of **operational methods**. Those are features, not bugs, in operational life. But they are disqualifying traits when the issue is **process selection** after a failure event—because process selection determines whether facts can be **compelled**, contradictions **tested**, and omissions **exposed**. A **Royal Commission** is not merely a longer meeting with better branding. It is a **truth-forcing machine: subpoena power, sworn evidence, adversarial testing**, and a **public record** wherever possible. A review, even an independent one, is structurally different. It tends to operate in **closed settings**, report in **summaries**, and leave the public unable to **verify** whether the hardest questions were fully confronted or carefully navigated around.

The Commissioner's intervention therefore lands on a precise democratic nerve: the distinction between **internal assurance** and **external accountability**. Internal assurance says, "we have reviewed ourselves through proper channels." External accountability says, "you do not get to be the final judge of your own adequacy after catastrophe." The legal and constitutional rationale for **Royal Commissions** exists because democracies recognise how power behaves under threat: it defaults to **containment**, not exposure; to **narrative management**, not maximal disclosure; to a preference for processes that are **controllable, scoped**, and less likely to produce rolling revelations that destabilise government and

agencies alike. None of this requires a conspiracy claim. It is the ordinary psychology of institutions.

The Prime Minister's reliance on **security-agency advice**, in that light, functions politically as both **shield** and **accelerant**. It shields by transferring authority: refusal is framed as **expert-endorsed prudence** rather than political calculation. Yet it accelerates **distrust** because it implicitly concedes that the **security apparatus** is not only an object of accountability but an active participant in determining accountability's boundaries. For a public already conditioned by repeated "prior interest / closure / catastrophe" sequences internationally, that is the precise configuration that produces the most corrosive inference: that the State is choosing a process that maximises **control** rather than **truth**.

The dispute matters, finally, because it reveals a deeper structural choice the country makes after each terror event. The choice is not simply "**Royal Commission** or not." It is whether national security becomes a domain where democratic government remains the **principal**, and agencies remain the **agents**—or whether, by habit and deference, the relationship begins to invert, with elected leaders citing the security state not as an input but as an authority that constrains what elected government will even attempt. When that inversion occurs, the tail does not merely wag the dog. The tail begins to define what the dog is allowed to know about itself.

ASIO HQ - Canberra





Chapter 2 — Bondi Beach as a National Fault Line

Bondi Beach is not simply a place; it is a national symbol. It is marketed to the world as open-air **Australianness**—sun, surf, freedom, proximity, and ease. It is also, in practice, a dense civic commons where tourists, locals, migrants, faith communities, and subcultures occupy the same public space without needing to agree on much of anything. That is why a mass-casualty attack at Bondi does more than produce grief. It turns an iconic site of everyday normality into a public question about whether the State can still guarantee the basic conditions of communal life: **safety, order**, and equal protection in the open.

On **14 December 2025**, that question became unavoidable. A **Hanukkah** gathering at Bondi—widely reported as “Chanukah by the Sea”—was attacked by two gunmen, with one killed at the scene and the other taken into custody after being wounded. Reporting describes **at least 15 deaths** and **dozens injured**, with the incident treated by authorities as a **terrorist attack**. The place itself amplified the shock. An attack on a secluded facility, a remote venue, or a secure building is one kind of trauma. An attack at **Bondi**, in open space, in daylight public imagination, is another. It communicates—whether intentionally or not—that no venue is too ordinary, no crowd too familiar, no celebration too integrated into civic life to be protected by default.

The immediate aftermath exposed three overlapping stress tests. The first was **social cohesion**. A Jewish community celebration attacked at one of the country’s most visible

public beaches forces the nation to decide whether it can speak honestly about rising **antisemitism**, while also resisting the spiral into collective blame or communal suspicion. The second was **counterterrorism readiness**. When the alleged attackers are described in public reporting as having been inspired by **Islamic State**, the question becomes not only how radicalisation occurs, but whether prevention systems—intelligence, policing, threat assessment, and disruption—are operating at the level the public has been led to believe. The third stress test was the most politically dangerous of all: **public confidence in state capacity**. A society can endure tragedy. What it struggles to endure is the suspicion that tragedy may have been preventable, that warnings were not acted on, or that institutional reflexes were shaped more by reputation management than by accountability.

This is where **Bondi** becomes a fault line, because what happens next determines whether the event is absorbed into the national story as a singular horror, or whether it becomes evidence of deeper institutional malfunction. A terrorism incident does not only test the perpetrator's capacity for violence; it tests the State's capacity for truth. That is why, within days, the debate moved rapidly from mourning into demands for structural explanation, including whether the attack could have been averted and whether government responses were proportionate to the severity of the failure alleged by many in the public sphere.

The Commonwealth's response set the terms for the trust battle. Prime Minister **Anthony Albanese** announced an independent review to assess whether the attack could have been prevented and whether there were gaps in legal and intelligence frameworks, while resisting calls for a federal **Royal Commission** on the basis that such inquiries take years. For families and critics, this structure raises a central concern: a review can deliver recommendations, but it rarely satisfies a public that suspects **blind spots**—missed signals, miscommunications, institutional silos—may have been decisive. When the public senses a gap between what it is told and what it suspects is true, the question is no longer merely “what happened?” It becomes “who gets to decide what we are allowed to know?”

Bondi therefore functions as a national fault line because it concentrates every modern tension into one event: **identity, security, terrorism, hate**, and the legitimacy of institutions that claim to protect the public while operating behind layers of secrecy. The beach is a symbol of the shared commons; the attack turned it into a referendum on whether the shared commons can still be defended, and whether the State can restore **public trust** without opening its decision-making to meaningful, independent scrutiny





Chapter 3 — The Warning Environment Before Chanukah

In the weeks and months before the **Bondi** massacre, the Jewish community in Australia was operating in a climate that many described not as episodic prejudice but as a sustained escalation in **antisemitic incidents**, ranging from street abuse and threats to vandalism and physical intimidation. The national incident picture was being tracked formally through annual reporting, with the Executive Council of Australian Jewry documenting **1,654** reported anti-Jewish incidents in its 2025 reporting cycle and emphasising that, even where year-on-year numbers fluctuated, the baseline had normalised into something “deeply troubling” in its human toll and everyday character. This was not simply an argument about sentiment. It was an argument about lived conditions: whether Jews could attend school, worship, or walk through public spaces without being treated as legitimate targets, and whether the State was calibrating its protective posture to match that reality.

Government responses to this environment had already been developing prior to Bondi. The Commonwealth appointed a **Special Envoy to Combat Antisemitism** in **July 2024**, and in 2025 released a public response outlining measures framed as institutional alignment against antisemitism—policy definitions, public service education, and targeted actions designed to reduce risk and improve capability across agencies. These steps mattered politically because they acknowledged that antisemitism was not merely a fringe phenomenon but a security and social-cohesion issue requiring structured intervention. Yet they also created a heightened expectation that, when credible warnings existed around specific events, protective measures would not remain abstract commitments but would translate into visible, practical safeguards.

Against that backdrop, the period leading into **Chanukah**—a time when public gatherings, ceremonies, and community celebrations become more frequent and more visible—was understood as a predictable stress point. This was not paranoia; it was risk logic. Public faith events concentrate people at known times and places, often with families and children present, and they can be readily targeted by offenders seeking symbolic impact. The environment is further sharpened when threat reporting and public rhetoric indicate that a group is being positioned, in the imagination of extremists, as a legitimate enemy. The question then becomes not whether a major event will be explicitly threatened in advance, but whether the overall threat environment is sufficiently elevated that prudent resourcing demands additional protection as a default.

This is the context in which reporting describes Jewish community security activity escalating from general vigilance into specific, written risk assessment in relation to the Bondi Chanukah gathering. ABC reported the existence of a leaked assessment by **Community Security Group NSW (CSG NSW)** dated **26 November 2025**, which warned that the event carried a **high risk**, including a heightened risk of **violent antisemitism**, alongside broader lone-actor threat considerations. The significance of that date is structural: it sits in the three-week window before the attack, at precisely the time when event planning, police engagement, and resourcing decisions would ordinarily be finalised or locked in. It also demonstrates that the warning environment was not merely retrospective storytelling. It was being articulated in advance, in a form that could—in principle—have been operationalised.

Alongside the risk assessment, reporting indicates that the Jewish community sought heightened protective presence for the event itself. 9News reported that requests were made for **additional officers** and that, after the attack, NSW Police declined detailed comment on operational questions on the basis that the incident is subject to a **critical incident investigation**, a **criminal investigation**, and a **coronial inquest**. The NSW Premier, **Chris Minns**, was also reported as acknowledging that CSG was “in contact” with police and that the police presence at the event was “clearly not enough,” language that—whatever the final legal and factual findings—reflects a political recognition that the operational posture did not match the perceived risk.

This is where the warning environment becomes more than an atmosphere; it becomes a contested chain of responsibility. Community security organisations are not police, and they do not control rosters, staffing levels, or the deployment of specialist capability. Their role is closer to threat-awareness, liaison, and risk communication. Yet when a community can point to an elevated climate of antisemitic hostility, a documented risk assessment dated weeks prior, and reported requests for increased protection, a powerful presumption forms in the public mind: that the system had enough information to treat the event as high risk, and that protective resourcing should have been correspondingly enhanced.

The State’s problem, in such circumstances, is not simply operational; it is institutional legitimacy. When public authorities cannot or will not disclose the decision chain—who received what, when they received it, how it was assessed, what options were available, what was requested, what was declined, and why—an information vacuum emerges. Into that vacuum flows the most corrosive force in democratic governance: the suspicion that warnings are treated as politically inconvenient until catastrophe makes them unavoidable. The gap between what the public can see and what it is told it cannot know becomes, in itself, a driver of instability, because it pushes the debate away from reform and toward distrust.



Chapter 4 — Community Risk Assessments and the Limits of “Informal” Warnings

Community security intelligence is a distinct species of warning. It is neither the product of statutory surveillance powers nor the output of classified collection. It is usually assembled from the ground level: incident reporting by community members, pattern recognition across repeat events, liaison conversations with local police, open-source monitoring of extremist rhetoric, and practical knowledge of how hate escalates in real public settings. Organisations such as **CSG NSW** describe their work in precisely those terms—protecting Jewish life through physical security support, training, CCTV, and “research,” built around a model that includes trained personnel and volunteers. In other words, community security groups sit in the seam between lived experience and formal policing. They are close enough to feel shifts in threat temperature early, but they are not endowed with the authority to deploy officers, compel compliance, or run the state’s operational machine.

The intelligence they generate tends to be practical rather than theoretical. It asks questions that rarely appear in national security briefs but matter enormously in the real world: what is the local incident trend, what symbolic date is approaching, what site is public-facing, what crowd profile is likely, what adversarial groups are active, what the risk of confrontation or lone-actor escalation looks like, and whether organisers should expect harassment or

violence. When that intelligence is reduced into a written assessment, it functions as a translation document—taking the community’s threat environment and rendering it into a form that authorities can act upon. That is why the reported existence of a leaked **CSG NSW** risk document dated **26 November 2025**—described by the ABC as warning that the Bondi “Chanukah by the Sea” event carried a **high risk**, including a heightened risk of **violent antisemitism** and broader lone-actor considerations—became so central to the post-attack debate. It represents, at least on its face, the precise mechanism by which community concerns are supposed to become operational decisions.

Yet the mere existence of a community risk assessment does not ensure it enters the state’s decision chain in a way that creates accountability. This is the first limit of “informal” warnings: the difference between a warning being produced and a warning being operationally received. If a document sits in the wrong inbox, is sent to a liaison contact without decision-making authority, is treated as advisory rather than actionable, or is not logged into a system that forces escalation, it may remain socially significant while being institutionally weightless. The ABC reported that a spokesperson for the NSW Police Minister said they were not aware of the document or any warning given to police. At the same time, reporting has quoted Premier **Chris Minns** saying that, based on media reports and his own inquiries, **CSG was in contact with NSW Police**, and that the police presence at the event was “clearly not enough.” Even without resolving that discrepancy, the structural point is clear: where warning flows are not formally captured, receipt itself becomes contestable, and contestability is a solvent that dissolves responsibility.

The second limit is that community intelligence can be treated as a “soft” input rather than a determinative trigger. Policing and security agencies are perpetually triaging competing risks—multiple events, finite personnel, unpredictable threats, and legal obligations that do not always align with community expectations. In that environment, community warnings can be discounted as “understandable anxiety,” particularly when agencies believe their own threat picture is more authoritative. The problem is that community groups often have superior visibility of certain kinds of threat escalation—especially harassment, intimidation, and hate-driven mobilisation—precisely because the community experiences it before it becomes a police statistic. When agencies interpret these warnings as subjective rather than intelligence-bearing, the system creates a built-in lag, where warnings become “real” only after violence supplies the proof.

The third limit is organisational fragmentation. Community warnings commonly move along relationship pathways rather than governance pathways. They are carried by liaison officers, event coordinators, local command contacts, or informal security networks. Relationship pathways are valuable for speed and trust, but they are weak in auditability. If the chain is not anchored to a clear procedural obligation—acknowledge receipt, log the warning, assess the risk against stated criteria, decide resourcing, communicate a decision back—then after an incident the story becomes a fog of recollection: who spoke to whom, what was “said,” whether it was “understood,” and whether it was “acted on.” That fog is precisely what drives public suspicion, because the public assumes—often rightly—that if a chain cannot be reconstructed, then the chain was not treated as important at the time.

The fourth limit is that even where a warning is received, it may not translate into visible protection due to operational or political constraints that remain hidden from the public. Some constraints are mundane: roster saturation, competing deployments, or a risk matrix that rates threats differently than communities do. Others are institutional: concerns about

setting a precedent, the optics of heavy policing at public religious events, or the fear of appearing to concede that a group is unsafe in the general community. After Bondi, 9News reported that NSW Police declined detailed comment on these operational issues, citing the matter as subject to a critical incident investigation, criminal investigation, and a coronial inquest. Silence may be justified procedurally, but it has a predictable social consequence: it leaves the public with only the inputs (warnings and requests) and the outcome (mass casualty), while the decision logic remains sealed. In that vacuum, distrust grows faster than any official statement can contain.

This is why the limits of “informal” warnings matter so much in a democracy. When a community flags heightened risk, and the state’s protective posture appears insufficient, the core issue is not whether the community’s warning was perfectly drafted or whether it used the correct bureaucratic channel. The core issue is whether the state has built systems that treat community intelligence as a legitimate early-warning sensor rather than as a public-relations irritant. If the warning chain depends on personal relationships and contested recollections, the system is structurally incapable of proving that it responded appropriately. And when the state cannot prove it responded appropriately, **public confidence** collapses into the presumption that it did not. That presumption is not a legal finding, but it is a political reality—and it is the exact condition that makes a coercive inquiry, capable of compelling documents and sworn evidence, not a luxury but a rational necessity for restoring trust.

Chapter 5 — Requests for Police Protection and Operational Triage

When a community asks police for heightened protection at a public event, the request is not a vague plea for reassurance; it is a request for a particular kind of operational posture. In a functioning security environment, such requests are supposed to trigger a structured process: liaison, assessment, planning, and resourcing. After the Bondi massacre, reporting indicated that Jewish community representatives had sought increased police coverage for the “Chanukah by the Sea” gathering, and the NSW Premier was quoted as saying that the police presence was “clearly not enough.” The phrase matters because it is not a moral statement; it is an operational verdict. It implies that the site was not protected at a level commensurate with the risk profile that, in hindsight, the community believed was obvious in advance.

Event policing, at its core, is a form of **operational triage**. Police commands never have unlimited personnel. On any given weekend, the same workforce is balancing domestic violence calls, road policing, custody management, patrol coverage, specialist deployments, and a rotating calendar of sport, nightlife, protests, festivals, and religious events. The decision to allocate extra officers to one location is always a decision to allocate fewer officers somewhere else. That is why requests for protection—especially from communities experiencing heightened hate—become a test of whether the state treats **targeted-group risk** as a central planning variable or as a discretionary add-on. The public may assume the default response is “protect the event,” but policing often experiences the request as “justify the uplift,” and the outcome turns on how risk is scored and how competing tasks are prioritised.

This is where the mechanics become decisive. Event policing is not only about the number of uniforms visible. It is about **where** officers are placed, **what** they are tasked to do, **what response time** can be achieved, and whether the posture is configured for deterrence,

disruption, or merely reaction. A high-risk posture typically requires more than a single patrol car driving past. It requires on-foot presence integrated into the crowd, clear command-and-control lines, rapid access to armed capability where relevant, and a pre-briefed understanding of evacuation routes, casualty management, and escalation triggers. Even small differences in posture can determine whether an attacker is deterred, interrupted early, or allowed to run through a crowd until a later stopping point. When people later say “there weren’t enough police,” they are often describing the felt absence of this layered posture—the difference between a symbolic presence and a system designed to defeat a fast-moving threat.

Operational triage also operates on timing. Most event security decisions are locked in before the event occurs, often days or weeks earlier, because rosters, logistics, and specialist support cannot be conjured at the last minute without destabilising the rest of the policing system. That is why the reported existence of a pre-event community risk assessment dated late November, and reported contact with police, is so important to accountability. If warning material is received early enough, it can change the resourcing plan. If it is received late, sent to the wrong level, or treated as advisory rather than actionable, the event runs on default settings. After the fact, the public tends to interpret default settings as neglect; agencies tend to interpret them as resource reality. A credible inquiry has to pierce this gap by tracing the actual timing: what was received, when, who made decisions, and what options were realistically available.

The phrase “not enough police on site” therefore has two meanings, and each points to different institutional failures. The first meaning is simple resourcing insufficiency: there were not enough officers physically present to create a deterrent footprint and to react quickly once violence began. The second meaning is strategic insufficiency: the presence that existed was not the right kind of presence—insufficiently positioned, insufficiently tasked, insufficiently prepared, or insufficiently connected to rapid escalation capability. In modern terror incidents, this distinction is not academic. The difference between stopping an attacker in the first moments and stopping an attacker after multiple magazine changes is measured in lives. In that sense, event policing is not merely about crowd management; it is about whether the state can project a credible guarantee that public faith gatherings will not be treated as soft targets.

The public debate is further inflamed when authorities do not provide operational clarity afterward. Reporting indicated NSW Police declined detailed comment on resourcing and decision-chain questions, referring to ongoing investigations and processes. That may be procedurally appropriate, but it has a predictable social consequence: the public sees only inputs and outcomes, not the logic that connects them. The inputs are the reported rise in antisemitic threat, community warnings, and requests for protection. The outcome is mass casualty on a public beach. In the absence of disclosed decision logic, the public fills the space with inference, and inference almost always collapses toward the harshest explanation: that the state knew, was warned, and did not act with seriousness.

This is the true political danger of operational triage. Triage is necessary, but triage is also morally combustible when it intersects with targeted communities. If a Jewish community flags heightened threat and seeks additional protection, and the policing posture later appears inadequate, the state cannot restore trust with general statements about “process” or “ongoing investigations.” It must be able to demonstrate that its triage system is rational, auditable, and defensible—that it has a coherent mechanism for translating credible community threat

signals into appropriate protective resourcing. Where it cannot demonstrate that, the argument for coercive inquiry strengthens. A Royal Commission is not demanded only because something went wrong; it is demanded because the public cannot see how triage decisions were made, and because “trust us” is no longer a credible substitute for a documented chain of accountability.



Chapter 6 — The Two Narratives: Antisemitic Targeting and ISIS-Inspired Ideology

In the aftermath of a terrorist attack, governments and agencies do not simply investigate what happened; they also decide how the event will be *understood*. That decision is rarely framed as “messaging,” but it functions as messaging all the same. The labels used in the first week—**what this was, who it targeted, why it happened, and what it means**—shape the public’s emotional map, the media’s organising logic, and the political system’s allocation of blame. In the Bondi case, two narratives emerged immediately and have operated in parallel: one emphasising **antisemitic targeting**, the other emphasising **ISIS-inspired ideology**. Each narrative is compatible with the other, but each pulls accountability in a different direction, changes which remedies feel “rational,” and influences what the public will remember years later.

The **antisemitic targeting** narrative is grounded in the victim profile and the setting: a Jewish Chanukah gathering at Bondi. Senior ministers have used direct language describing the event as a targeted attack on Australia’s Jewish community. In a joint press conference, the Attorney-General described what occurred at Bondi as a “targeted, antisemitic terrorist attack

on the first day of Chanukah.” This framing is about the *object* of the violence: a community attacked for who it is. Its moral logic is straightforward, and that simplicity is politically powerful. If the attack is framed primarily as **antisemitism**, then the questions that follow are about hate and protection: whether antisemitic threats were escalating, whether community warnings were heeded, whether police resourcing was sufficient, and whether government responses to antisemitism had been performative rather than operational.

The **ISIS-inspired ideology** narrative focuses on the *driver* of the violence: radicalisation pathways and terrorist inspiration. The Prime Minister has repeatedly used language describing the attack as “ISIS inspired,” calling it an “ISIS inspired atrocity,” while still acknowledging solidarity with the Jewish community and describing the incident as terrorism. Reuters reporting likewise describes officials saying the alleged attackers were inspired by Islamic State, in the context of an attack at a Jewish Hanukkah celebration. This framing is about the machinery of contemporary terrorism: extremist content ecosystems, online networks, watchlist thresholds, disruption powers, and inter-agency intelligence sharing. It naturally pushes the policy response toward counterterrorism architecture—reviews of powers, legal settings, and intelligence frameworks—rather than toward the narrower operational question of whether a specific Jewish event was protected at an adequate level.

This is where framing begins to move the locus of accountability. When the attack is framed primarily as **antisemitism**, responsibility concentrates around the protective obligations of the State to a targeted minority. The public asks whether visible threat escalation was dismissed, whether pre-event risk assessments were treated as noise, and whether police triage downgraded the safety of a Jewish gathering in a heightened hate environment. When the attack is framed primarily as **ISIS-inspired terrorism**, responsibility is redistributed into the more ambiguous terrain of intelligence work: what constitutes actionable intelligence, what can be shared, what cannot be monitored indefinitely, and whether prevention is possible without a level of intrusion society will not tolerate. The effect is not that antisemitism disappears; it is that the argument becomes harder to resolve in public, because it is forced into classified spaces and contested thresholds that governments can describe as inherently uncertain.

Governments often claim these frames are simply different facets of the same truth, and in one sense that is correct. A person can be motivated by **ISIS ideology** and select Jews as the target because that ideology defines Jews as the enemy. Yet in public life, sequencing and emphasis matter. If the public repeatedly hears “ISIS-inspired” as the headline descriptor, the political mind begins to treat the attack as primarily a counterterrorism failure that requires more powers, more surveillance, and more secrecy. If the public repeatedly hears “antisemitic terrorist attack” as the headline descriptor, it begins to treat the event as a failure of targeted-community protection that requires more visible policing, more accountability around pre-event warnings, and a more aggressive response to hate and intimidation. The same facts, emphasised differently, can produce different “obvious” remedies, and the chosen remedies then retroactively justify the chosen framing.

This tension is visible in the Commonwealth’s institutional response. Albanese announced an independent review to examine whether the attack could have been prevented and whether agencies and legal frameworks were adequate, while resisting calls for a federal **Royal Commission** on the basis that such inquiries take years. A review process of this kind fits naturally within the **ISIS-inspired** frame: it centres intelligence and law enforcement

mechanics, gaps in powers, and coordination settings. Families demanding a Royal Commission, by contrast, are signalling that the public is not satisfied with a technocratic counterterrorism audit; it wants a process capable of compelling the entire decision chain, including what was known, what was warned, what was requested, what was resourced, and why.

Framing also shapes **public memory**, and public memory determines whether a society learns or simply mourns. The struggle over whether to name antisemitism explicitly can be seen in symbolic decisions after Bondi. Reporting described pressure for public commemorations to acknowledge the antisemitic nature of the attack rather than relying on generic unity imagery, with changes made after criticism that specificity was being avoided. This is not a superficial dispute. If the attack enters national memory as “a terror incident,” the targeted character of the violence can dissolve into abstraction. If it enters memory as “an antisemitic terror attack at a Jewish celebration,” it becomes harder for institutions to retreat into generalities, because the identity of the target group remains embedded in the public record.

In practice, the most consequential issue is not which narrative is “true,” but whether framing choices are being used—consciously or unconsciously—to steer scrutiny away from certain questions. The **antisemitism** frame forces uncomfortable operational questions about event protection, policing triage, and the handling of community warnings. The **ISIS-inspired** frame forces uncomfortable questions about intelligence thresholds, watchlist logic, and prevention capability. A democratic society cannot afford to treat these as competing political tribes. It needs a factual process that can hold both frames in view and test where the real institutional **blind spots** were, because the deeper risk is not merely the next attack—it is the corrosion of **public trust** created when framing substitutes for disclosure and when the public senses that language is being used to manage accountability rather than reveal it.



Chapter 7 – ASIO Prior Interest and the Closure Problem

One of the most combustible facts in the Bondi record is not a rumour but a documented public admission: **ASIO examined Naveed Akram in 2019** in relation to alleged connections to an **Islamic State** network, and that attention was later **closed** after ASIO concluded there was no evidence he had been radicalised or posed an **ongoing threat**. After the attack, this detail has functioned like a match thrown into dry grass, because it collides with the public’s intuitive model of intelligence work. Ordinary citizens assume that once an individual has crossed into intelligence interest—particularly where extremist links are alleged—some kind of durable protective perimeter follows. The reality is more austere: intelligence services operate under triage, thresholds, and finite capacity, and the routine administrative act of **closing** a line of inquiry is a structural necessity of the system, even though it can look like negligence when catastrophe follows.

This is the **closure problem**. The state can legitimately say that it cannot keep everyone under perpetual surveillance; the public can legitimately reply that this is precisely why closure decisions must be exceptionally disciplined, documented, and reviewable. ABC reporting after Bondi described the “grim reality” that the volume of threats makes persistent surveillance impossible and emphasised that little is publicly known about the alleged offenders’ activities between 2019 and the 2025 attack—an information gap that inflames suspicion precisely because it sits in the space where intelligence agencies claim expertise. The public hears “no ongoing threat,” yet also hears that a mass-casualty attack occurred at a Jewish religious gathering, with federal authorities treating the incident as terrorism. The

question that follows is predictable: if an earlier investigation existed, what changed, what was missed, and who would have known?

The closure problem is not solved by repeating that a case was closed. It is solved only by reconstructing the decision logic that produced closure and then testing whether that logic remained rational as circumstances evolved. This is especially so when the attack itself carries multiple attributes that, at least to the lay observer, read as “obvious escalation”: a high-profile Jewish public event, a documented community warning environment, and alleged access to firearms through a licensed household. None of that proves that ASIO could have prevented the attack, and none of it proves that ASIO acted improperly in 2019. But it does underscore the stakes: closure is not merely an administrative endpoint; it is a risk decision with downstream consequences.

Bondi also illustrates the practical limits of how intelligence interest translates into operational prevention. Even where an agency has information, the system is segmented. Intelligence interest does not automatically map onto policing deployment at a specific event, and intelligence agencies do not roster patrols. Where the warning chain is fractured—community assessments on one side, event-policing triage on another, intelligence triage in a separate compartment—the system can produce a paradox: each component can claim it acted rationally within its remit, while the combined system fails to protect the public. The Reuters account of the Commonwealth response captures the political essence of this dilemma: Albanese has announced an independent review to assess whether the attack could have been averted and whether there were gaps in legal and intelligence frameworks, while resisting calls for a Royal Commission on the basis that it would take years. That response implicitly concedes the possibility of systemic gaps while seeking a faster, more contained pathway to reform.

The deeper significance of the closure problem is that it recurs internationally, and therefore it cannot be dismissed as an isolated embarrassment. Across multiple jurisdictions, a familiar pattern appears: a future attacker was **previously examined**, later **deprioritised**, and only then re-emerges as a perpetrator. The pattern does not prove malice or intent; it more often reflects the harsh arithmetic of risk triage. But precisely because the pattern is recurring, the public expectation that closure decisions should be **auditable**, and subject to independent scrutiny after catastrophe, is not unreasonable—it is the minimal condition of democratic legitimacy when agencies operate behind secrecy.

That is why Bondi cannot be contained as a question of whether ASIO made the “right call” in 2019 in the abstract. The only question that matters is whether the system’s triage and closure logic remained defensible against the evolving context, and whether any subsequent indicators should have triggered a re-opened posture, escalated sharing, or different protective settings. The government’s decision to commission the Richardson review explicitly acknowledges that these are the questions now at the centre of legitimacy: what agencies knew, what they did, what they could have done, and whether the attack was preventable. In a trust-fractured environment, however, closure logic cannot be restored by assurances alone. It requires a process capable of compelling documentary truth about how closure was reached, how risk was rescored over time, and where the system’s blind spots became fatal.



Chapter 8 — International Precedent: Orlando, Boston, 7/7, Manchester

The Bondi debate did not emerge in a vacuum. Democracies have repeatedly faced a destabilising sequence in which a future attacker is **previously scrutinised** by the security apparatus, later **deprioritised** or formally **closed**, and then reappears as a perpetrator. This pattern is not proof of intent or complicity. It is, however, a recurring institutional stress test that exposes the same structural **blind spots** across jurisdictions: how intelligence agencies triage uncertainty, how warnings are translated into operational action, and how “closure” decisions can look rational on paper but catastrophic in hindsight.

In the United States, the **Pulse nightclub** attack in Orlando became a landmark example of the closure problem in public consciousness. Mainstream reporting and post-attack briefings described how the FBI had previously scrutinised **Omar Mateen** in earlier investigations before concluding there was insufficient basis to continue and closing the inquiries. The significance of that sequence is not that the FBI “knew” an attack would occur; it is that the public learns, after bodies are counted, that the perpetrator had already passed through the system’s filters. That discovery predictably produces a legitimacy shock, because the public does not experience “closure” as an administrative necessity; it experiences it as a failure of foresight. The deeper institutional lesson is that the decisive questions are rarely “Was he investigated?” but rather “What indicators were weighted, what thresholds were applied, what was deemed actionable, and what would have had to change for the assessment to remain open?”

The **Boston Marathon** bombing illustrates the same architecture of vulnerability, but with a different trigger: foreign intelligence warning and bureaucratic process. The joint inspectors general report into information handling prior to the April 2013 bombing records that **Tamerlan Tsarnaev** came to the FBI's attention in 2011 based on information from Russia's security service alleging radical Islamist adherence and potential travel to join underground groups, and that the FBI-led JTTF conducted an assessment and **closed it three months later** after finding no "nexus" to terrorism. This kind of material is especially instructive because it shows how institutional blind spots can be produced without malice: warnings arrive, an assessment is conducted, boxes are ticked, and the case is closed because the available information does not cross internal thresholds. The failure mode is not necessarily laziness; it is that the system is designed to make closure decisions under conditions where the most dangerous intent may be latent, unspoken, or temporarily dormant.

The United Kingdom's 7/7 experience places the same problem inside an environment explicitly described by oversight bodies as one of **selectivity and risk-bearing** due to finite resources. The Intelligence and Security Committee's report on the 7 July 2005 attacks sets out that MI5 must be selective, that targets move between investigative tiers as intelligence changes, and that decisions are constrained by proportionality and resource allocation. The importance of that framing is that it normalises, in official language, the central reality the public often resists: intelligence agencies routinely accept a degree of risk because they cannot cover everything. Yet the report also acknowledges the "limits of coverage" and the possibility that attacks can be planned without detection. The recurring blind spot revealed in such post-event reviews is not that agencies had no information; it is that information often sits in fragments—tiered, compartmented, and deprioritised—until a coherent pattern becomes visible only after the attack has supplied the missing context.

The **Manchester Arena** case is perhaps the clearest illustration of how public inquiries crystallise institutional lessons into findings that damage legitimacy. The Chair of the Inquiry made explicit that the purpose of the exercise was to examine whether more could and should have been done to stop the bomber, and that criticism would be unavoidable where identifying improvements required it. Subsequent reporting on the Inquiry's conclusions stated that MI5 missed a significant chance to prevent the attack and that the bomber had previously been a subject of interest before his case was closed. The Manchester lesson is not merely that an opportunity was missed; it is that large-scale counterterrorism systems can fail not only at the level of "collection" but at the level of **internal communication**, escalation, and timely action. Where one officer does not share or elevate intelligence quickly, the machinery can behave as if it never possessed the information at all.

Across these precedents, a set of institutional blind spots appears with relentless consistency. First, there is the blind spot of **threshold governance**: agencies must decide what crosses from "concerning" to "actionable," and those thresholds are shaped by law, resourcing, policy, and culture. Second, there is the blind spot of **administrative closure**: once a case is closed, systems tend to treat it as resolved unless fresh triggers force re-opening, even though real-world radicalisation and intent can be nonlinear, episodic, and opportunistic. Third, there is the blind spot of **fragmented warning flows**: foreign intelligence, community reporting, policing observations, and online signals often arrive through different channels that do not integrate cleanly, producing a situation where no single node feels responsible for the whole risk picture. Fourth, there is the blind spot of **posture translation**: even where intelligence exists, converting it into operational protection—at a venue, on a date, with sufficient capability—requires a chain of decisions that can fail at multiple points without any actor

believing they have acted irrationally within their remit. Fifth, and most corrosive, there is the blind spot of **accountability opacity**: when the public cannot see how triage decisions were made, it assumes the worst, and the legitimacy of institutions begins to depend not on effectiveness but on whether they can compel belief without disclosure.

Bondi sits inside this international pattern, which is precisely why “prior interest” cannot be dismissed as an irrelevant historical detail. The recurring lesson from Orlando, Boston, 7/7, and Manchester is that democratic societies repeatedly discover, after catastrophe, that the attacker had previously intersected with the system—sometimes through investigations, sometimes through warnings, sometimes through fragments of intelligence that never coalesced into prevention. The institutional question is therefore not whether such intersections are scandalous in themselves, but whether the system is designed to prevent the predictable failure modes that these precedents keep revealing: closure without robust revisit logic, triage without auditable justification, information without escalation discipline, and protective posture that does not match the risk environment until after the event has proved the risk was real.

Chapter 9 — Firearms Access Pathways and Household Risk

Modern firearms regulation is built on a reassuring premise: if the state licenses the “right” people, and those people comply with storage and handling rules, firearms can exist in civilian hands without producing intolerable risk. The Bondi case exposes the fragility of that premise, not because it proves malice by regulators, but because it highlights a structural reality that licensing systems often underweight: **household access pathways**. A person does not need to hold a licence to be endangered by the presence of guns; they need only to be able to reach them. That is why, when public reporting indicates the father held a valid firearms licence while the son later became the alleged perpetrator, the question that follows is not rhetorical. It is operational: what safeguards existed to prevent access by a non-licensed household member, particularly one who had previously intersected with counterterrorism concern?

The public’s intuitive position is simple. If the state knows that an individual has been under intelligence attention for alleged extremist links—whether or not a threat was ultimately confirmed—then firearms inside that individual’s living environment appear, at minimum, as a foreseeable risk vector. The legal system, however, tends to treat risk in segmented compartments. Licensing decisions are often made on the basis of the licensee’s record, suitability, and compliance history, rather than a more expansive inquiry into household dynamics, vulnerable access points, or the presence of high-risk persons living in the same premises. In other words, the system can be rigorous on paper about the applicant while remaining structurally weak on the reality that guns are stored somewhere, and other people exist in that environment.

This is where the logic of “lawful possession” becomes insufficient. Safe storage rules matter, but they are not a guarantee against access. Households contain patterns of routine, trust, and familiarity that can defeat procedural safeguards. Keys are discovered, combinations are learned, storage practices drift, and complacency accumulates. It is easy for governments to speak as if storage compliance is a binary—compliant or non-compliant—

when in reality storage is a spectrum of behaviours under daily stress. The problem is not simply “bad people breaking rules.” It is that predictable human factors—fatigue, routine shortcuts, assumptions that “nothing will happen”—create the very access pathways that attackers exploit. In a scenario where a household member is motivated, determined, and intent on violence, “lawful licensing” does not address the core question: whether the system had mechanisms to anticipate and neutralise household risk.

In the Bondi case, the household risk question intersects directly with the closure problem discussed earlier. Public reporting indicates that the son, **Naveed Akram**, previously came to ASIO’s attention years earlier and was assessed as posing no ongoing threat after investigation. Even if that assessment was reasonable at the time, the presence of firearms within his living environment would be understood by any ordinary risk manager as an aggravating factor. It converts ideological volatility—whether latent or emerging—into immediate practical lethality. That does not mean ASIO should have taken control of a household’s firearm storage. It means the broader system should have been structured so that prior extremist concern could be reflected in firearms licensing and household suitability settings. The public’s outrage is not irrational; it is the predictable reaction of people who assume that government systems talk to each other, when, in reality, those systems often do not.

It is telling that, in response to Bondi, NSW moved quickly to tighten firearms law in ways that explicitly acknowledge the household pathway risk. Reporting describes new provisions intended to prevent granting a firearms licence not only to someone investigated for terrorism-related matters but also to someone who lives with a person who has been investigated. This is a quiet admission of a prior blind spot: that a household is a risk unit, not merely an individual applicant. The state’s previous model, like many jurisdictions, implicitly treated the licence holder as the relevant object of control. The revised model recognises that the relevant object of control is often the **environment of access**.

The deeper point is that firearms licensing regimes are not designed primarily for counterterrorism. They are typically designed for crime prevention, safety, and suitability in ordinary settings. But terrorist events exploit precisely the gaps created by ordinary governance models. A system built for incremental, everyday misuse can be outpaced by offenders seeking mass casualty. That is why “household access” becomes critical. It sits in the overlap between different worlds: intelligence assessments, policing triage, firearms regulation, and the lived reality of family relationships. Because it sits in that overlap, it is also a prime location for institutional buck-passing. Intelligence agencies can say licensing is not their remit; licensing authorities can say intelligence is not disclosed to them; police can say they are constrained by process; and government can say every system worked as designed. Meanwhile, the public sees the plain practical outcome: if firearms were in the household, access existed, and access was exploited, then “as designed” was not good enough.

The argument here is not that lawful gun owners are culpable for terrorism. It is that governments have repeatedly sold the public a simplified story: that licensing equals safety, and that compliance equals control. The Bondi case, and the reforms that followed, point to a more uncomfortable truth: licensing is necessary but not sufficient, because it does not automatically manage household access pathways. In high-risk contexts—especially where prior extremist concern exists—those pathways are not hypothetical. They are foreseeable. And when foreseeable risks become realised tragedies, trust collapses unless the state can

demonstrate that it had a credible mechanism to identify household risk and reduce it before catastrophe occurred.



Chapter 10 — Travel, Networks, and the Investigative Fog of “Red Flags”

In the first phase after a mass-casualty terrorist event, public attention often fixates on “red flags” as if they were self-explanatory proof. Travel patterns, social contacts, and ideological associations become symbols that people treat as settled evidence of a hidden pipeline. The problem is that “red flags” operate in an evidentiary fog. They are often strong **investigative triggers** while remaining weak as **public proof**. Bondi illustrates this tension sharply because a single fact—reported travel by the alleged offenders to **Davao City** in the southern Philippines in the weeks before the attack—can be simultaneously meaningful, suspicious, and still inconclusive without a documented chain of corroboration.

Public reporting indicates that Australian and Philippine authorities focused quickly on the pair’s November 2025 travel, treating it as a central investigative line of enquiry rather than a peripheral curiosity. ABC reporting describes police scrutiny of the men’s movements in the Philippines and the attention given to identifying locations, tracing activity, and clarifying the purpose of their time in Davao. SBS reporting similarly notes that the trip became a focus of counter-terrorism investigation while emphasising that what authorities have not confirmed—and may never fully uncover—is what the men did while in the Philippines. This is the correct investigative posture: travel shortly before an attack is neither “proof” nor

“irrelevant.” It is a lead that must be tested against records, timelines, and corroborated behaviour.

The reason the Davao trip became so combustible in public debate is that “southern Philippines” is often perceived—fairly or unfairly—as synonymous with insurgency, militant networks, and jihadist facilitation. That perception is not wholly baseless; Mindanao has a long history of militant activity and has been linked in past years to ISIS-aligned violence. Yet perception is not the same as operational reality, and post-event narratives can quickly outrun what the evidence will bear. Reuters reported that the Philippine National Security Adviser stated there was **no evidence** the suspects received military or terrorist training during their visit and that immigration records indicated they stayed in Davao from **1 to 28 November**, spending most of their time secluded in a budget hotel. Al Jazeera likewise reported Philippine officials saying there was “no evidence” of training, underscoring the key point: travel to a region associated in the public mind with extremism does not, by itself, establish contact with extremist networks or the receipt of operational training.


This is where the discipline of inference matters. A travel “red flag” can support several competing hypotheses, and responsible analysis must hold them apart rather than collapsing them into a single insinuation. The most alarming hypothesis is structured facilitation: travel for training, financing, or meeting an organised network. A second hypothesis is ideological reinforcement: travel that is meaningful to identity and commitment, but not operationally directed. A third is strategic misdirection: travel that looks like a red flag precisely because it is meant to draw attention away from other planning that occurred elsewhere. A fourth is the least dramatic but often the most common: travel is simply travel, and its relevance is limited to timing rather than causation. In Bondi, what makes the travel fact significant is not that the public can prove which hypothesis is correct, but that the **state must** be able to prove which hypothesis is false or unsupported, and why, using records that can withstand scrutiny.

The fog thickens further when “networks” become an elastic term. A person can be described as linked to extremist networks based on online consumption, ideological sympathy, follower relationships, or proximity to a controversial preacher, without ever having a direct operational connection to a terrorist organisation. The ABC reported that at least one of the alleged gunmen was known to authorities as a follower of a notorious Sydney pro-Islamic State cleric and that investigators were probing how the pair spent time in the Philippines. The Washington Post reported that Australian police were investigating whether the trip to the Philippines involved links to terrorist groups. These are legitimate investigatory angles, but they underscore a central problem for public trust: “linked to” can mean anything from ideological consumption to real-world facilitation, and without compelled disclosure the public cannot know which sense is intended.

A further complication is that, even when authorities have records—hotel receipts, CCTV, immigration logs—those records rarely produce the kind of clean narrative the public wants. ABC reporting describes investigators examining footage and hotel recollections as part of tracing the men’s movements in Davao. Reuters reported that the pair reportedly remained largely secluded and had little external activity, a pattern that can be interpreted in multiple ways: as evidence against training claims, as evidence of covert planning, or as mundane behaviour inflated into significance by the context of what later occurred. This is why “red flags” are so dangerous to public reasoning. They feel like proof because they appear to explain the horror, but they often remain ambiguous until integrated into a larger chronology of preparations, communications, procurement, and intent.

Bondi's travel dimension also intersects with the institutional question of how "red flags" are meant to trigger state action before tragedy. Reuters has reported that the alleged offenders were inspired by Islamic State and that the independent review announced by Albanese will examine whether the attack could have been averted and whether gaps existed in legal and intelligence frameworks. That inquiry question is not abstract. If travel to a high-sensitivity region occurred shortly before the attack, the critical accountability issue is whether any part of the system—border settings, intelligence triage, joint counterterrorism frameworks—treated the travel as an elevated risk indicator, and if not, why not. Conversely, if agencies did treat it as a lead but could not lawfully act, then the policy question becomes whether the legal and operational thresholds are fit for purpose or whether the public has been sold an inflated understanding of what "counterterrorism" can do in real time.

The responsible conclusion, at this stage, is not that travel proves training, and not that travel proves innocence. It is that travel is a powerful investigative trigger that must be tested against documentary proof, and that the public cannot be expected to restore trust while being asked to accept a narrative without seeing the chain of evidence that supports it. In a case where a beachside religious celebration became a mass-casualty scene, "red flags" are not enough. They are precisely the kind of contested indicator that requires a process capable of compelling records, examining decision chains, and distinguishing what was merely suspicious from what was provable—because the difference between suspicion and proof is the difference between commentary and accountability.



Chapter 11 — Information-Sharing Failures: What Moves, What Sticks, What Dies in Silos

Large security systems rarely fail because no one knew anything. They fail because **knowledge does not move** in the way prevention requires. Bondi has become a live case study in this dynamic, not because the public currently possesses the full documentary chain, but because the publicly reported inputs point to multiple systems that *should* have intersected—community security assessment, event policing, counterterrorism intelligence, and firearms licensing—while the outcome suggests that, if intersection occurred, it did not produce a protective posture capable of stopping mass murder. The core question is therefore not merely whether information existed, but **what information moved**, what was absorbed and acted upon, what was stalled, and what effectively died inside institutional silos.

In theory, Australia’s counterterrorism architecture is built around coordination. Intelligence agencies collect and assess; police disrupt and deploy; licensing and regulatory systems reduce access to dangerous capability; and government provides policy and oversight. In practice, each node has different legal authorities, different secrecy regimes, different incentives, and different definitions of what counts as “actionable.” This produces a predictable failure mode: each node holds fragments that are meaningful in combination but unremarkable in isolation. When those fragments are not integrated, the system behaves as if it never possessed them at all.

Bondi presents a straightforward illustration of the fragmentation problem. Reporting indicates a Jewish community security organisation produced a written risk assessment dated **26 November 2025** describing the Chanukah gathering as high risk, including heightened risk of **violent antisemitism**. Reporting also indicates the community sought increased police protection for the event, and the NSW Premier publicly stated that police presence was “clearly not enough.” Reporting further indicates that the alleged perpetrator **came to ASIO’s attention in 2019**, was investigated for a period, and was later assessed as posing no ongoing threat at that time. Finally, the post-attack debate has centred on household firearms access pathways and the subsequent tightening of NSW firearms laws to address terrorism-related risk, including household proximity. These are not trivial signals. They are precisely the kind of signals that, if combined, would have justified a substantially hardened protective posture. The fact that the system did not appear, from the outside, to behave as if it had that combined picture is the essence of the information-sharing question.

The mechanics of information-sharing failures are more banal than conspiracy and more corrosive than incompetence. One common mechanism is **classification and compartmentalisation**. Intelligence is often held at a classification level that makes it difficult to share widely, not merely because agencies are protective, but because disclosure risks sources, methods, and international relationships. Even within government, counterterrorism information may be shared only with those who have a “need to know,” a concept that protects secrets but can also prevent prevention when the “need” is interpreted narrowly. This creates a paradox: the more sensitive information is, the more likely it is to be trapped where it cannot trigger operational change. When the public later asks why the system did not act, officials can respond that they cannot discuss what was known—an answer that may be true but that accelerates distrust because secrecy is indistinguishable from evasion in the public mind.

A second mechanism is **threshold mismatch**. Intelligence agencies work in probabilities and patterns; licensing and regulatory systems work in eligibility criteria; operational policing works in deployment capability and immediate threat. An intelligence assessment that someone is “of interest” may not meet the threshold for action in licensing, particularly if the relevant legal standard is framed around convictions, charges, or demonstrable threat. Conversely, licensing may legally permit possession because the licensee appears suitable, while intelligence may regard the household environment as a latent risk. In such cases, the system can be internally consistent while being externally dangerous: each node makes decisions that are defensible within its own rules, yet the combination of those decisions produces an avoidable access pathway. The speed with which NSW moved to restrict firearms licensing linked to terrorism investigation, including household proximity, suggests a recognition that prior settings were not adequately aligned to risk realities.

A third mechanism is **organisational self-protection**, which does not require bad faith to be harmful. Agencies and units are incentivised to avoid creating discoverable records that could later be used to assign blame. “Informal” liaison can be preferred to formal written escalation precisely because it is flexible and deniable. The result is a common post-incident pathology: disputes about who was warned, what was said, whether something was “received,” and whether a request was “made” in the right format. The ABC’s reporting that a ministerial spokesperson said they were not aware of the community risk document sits alongside other reporting describing community-police contact and political acknowledgment that on-site presence was inadequate. This is not necessarily evidence of deception; it is evidence of a warning ecosystem that lacks robust, auditable hand-off points.

A fourth mechanism is **resource competition**. Even where information moves, it must still compete for attention and personnel. Police commands triage events; counterterror units triage subjects; licensing authorities triage applications. In each context, information that does not demand immediate action can be buffered into “watch and wait.” Over time, “watch and wait” becomes “closed,” and “closed” becomes “not a current issue.” The closure problem, once embedded, becomes self-reinforcing: a case that is closed is assumed to be low risk; therefore fewer resources are allocated; therefore fewer new indicators are collected; therefore closure appears justified. That loop is rational within the system but can be fatal when a previously low-priority subject rapidly escalates to violence.

A fifth mechanism is **jurisdictional fragmentation**. Counterterrorism intersects Commonwealth and state responsibilities, while firearms licensing is primarily state-based, and community security liaison may operate at local-command level. The more nodes involved, the easier it becomes for each to assume the “other side” has the problem in hand. After Bondi, the Commonwealth established a review under Dennis Richardson to examine whether the attack could have been averted and whether there were gaps in frameworks and information sharing. That decision is an implicit admission that the key questions are systemic rather than purely local: what moved between agencies, what did not, and why.

The deepest issue is that information-sharing is not primarily a technical problem; it is a governance problem. If the system is designed so that the default posture is compartmentalisation, and the burden sits on someone to justify escalation, then prevention depends on individual judgment under stress. If, instead, the system is designed with mandatory escalation triggers—documented receipt of community high-risk assessments, mandatory threat review for events flagged as high-risk, formalised pathways from

counterterrorism assessment to licensing and household risk screening—then prevention becomes less dependent on individual heroics and more on structural reliability.

Bondi has made one point unavoidable: a security state that cannot demonstrate how information moves cannot sustain public trust after catastrophe. People will accept that perfection is impossible. What they will not accept is an architecture where warnings can be issued, requests can be made, and prior intelligence interest can exist, yet the system remains incapable—at least in the public’s view—of connecting those facts into a protective posture. The public does not demand omniscience. It demands an accountable chain: what was known, who knew it, who shared it, what was done with it, and why. Until that chain is exposed and tested, the most dangerous thing in any modern democracy will continue to grow in the shadow of every terror event: not only fear of violence, but fear that institutions are no longer structurally capable of telling the truth about their own failures.



Chapter 12 — The “Let-It-Happen” Hypothesis and the Standard of Proof

When a terror event exposes multiple warning indicators, a certain hypothesis reliably appears in the public mind. It is not always spoken in polite company, but it shapes the emotional logic of the debate: the possibility that agencies did not merely fail, but **allowed** the event to occur—whether to justify expanded powers, to shift public opinion, to protect reputations, or to advance a strategic agenda. This is the “**let-it-happen**” hypothesis, and it is often described in Machiavellian terms: that the security state may tolerate controlled catastrophe because catastrophe is politically productive.

Bondi, given the reported combination of warning environment, requests for protection, prior intelligence interest years earlier, travel red flags, and firearms access pathways, is exactly the kind of case in which this hypothesis becomes socially contagious. Yet a serious book cannot treat contagion as evidence. The hypothesis must be handled as precisely what it is: a **possible claim** that requires an unusually high standard of proof before it can be treated as credible. The purpose of setting that standard is not to protect agencies from criticism. It is to protect the public from confusing suspicion with truth, and to protect legitimate accountability from being discredited by overreach.

The Machiavellian thesis can be stated narrowly. It does not require that agencies “ordered” the attack. It requires that at least one relevant authority had **actionable foreknowledge** and made a deliberate choice to **stand down**, to **avoid disruption**, or to **permit escalation** because the perceived downstream benefits outweighed the costs. That is a specific claim about intent. And claims about intent—especially institutional intent—require evidence that is qualitatively different from evidence of negligence, misjudgment, or bureaucratic failure.

To treat a “let-it-happen” thesis as credible, the evidentiary burden would typically need to include one or more of the following elements, not as conjecture but as demonstrable fact. First, documentary proof of a **stand-down** or non-intervention decision that is not adequately explained by lawful triage or resource constraints—emails, minutes, tasking directives, or recorded instructions showing that disruption was deliberately avoided for reasons unrelated to capability or lawfulness. Second, evidence that **actionable intelligence was suppressed**—not merely missed, not merely delayed, but actively withheld from those who would have used it to prevent the attack. Third, proof of an informant or controlled asset with **foreknowledge** whose handlers permitted escalation rather than treating the information as a prevention trigger. Fourth, evidence of deliberate obstruction of protective measures—such as blocking resourcing after specific warnings had been received and understood, or refusing reinforcement in circumstances where reinforcement was both feasible and ordinarily mandated. Fifth, a pattern of conduct that cannot be plausibly explained by normal failure modes: repeated decisions, repeated non-actions, repeated misclassifications, or repeated compartmentalisation that collectively point to intention rather than accident.

Absent evidence in that class, the most defensible conclusion is not Machiavellian enabling but **systemic failure**. Systemic failure can be severe. It can include incompetent risk scoring, siloed information, bureaucratic denial, reputational defensiveness, and operational triage that undervalues targeted-community risk. But it remains distinct from intentional permissiveness. The distinction matters because the remedies are different. Systemic failure is repaired through governance, auditability, and operational reform. Intentional permissiveness—if ever proven—would represent a collapse of constitutional legitimacy requiring an entirely different response.

The problem is that the evidence required to distinguish these categories is rarely visible without compulsion, and rarely visible quickly even with it. Intelligence and counterterrorism processes operate behind multiple layers of secrecy. The most relevant material is often held in classified compartments and framed as source-sensitive, meaning that even the existence of certain intelligence cannot be admitted publicly. Decisions are frequently recorded in language designed to preserve defensibility: “insufficient grounds,” “no nexus,” “low priority,” “resource constraints,” “no ongoing threat.” Such phrases may be true, but they also create **plausible deniability** if a decision is later questioned. Moreover, the most consequential failures may not be written as decisions at all. They can be omissions: a

warning not escalated, a liaison note not entered, a briefing not sent, a request not logged. Omissions leave little forensic trace, which is why they breed conspiracy narratives: the absence of evidence becomes interpreted as evidence of concealment.

There is also a statistical trap that makes “let-it-happen” narratives psychologically persuasive. Security agencies investigate large numbers of people. Some of those people will later commit acts of violence, whether or not the agency acted properly. This produces an intuitive “they knew” conclusion in the public mind, even where the reality may have been a rational triage decision under uncertainty. The trap is intensified by the closure problem: a case closed years earlier appears, in hindsight, like a decision that should have been revisited, even if there was no contemporaneous trigger visible to the system at the time. In this environment, the public oscillates between two extremes: blind trust and cynical certainty. A responsible inquiry must resist both.

Bondi therefore demands a disciplined separation between hypothesis and proof. It is legitimate to say that the combination of reported indicators creates conditions where the public will naturally suspect deeper institutional culpability. It is legitimate to say that such suspicions cannot be resolved by reassurance, because reassurance is not evidence. It is legitimate to argue that only a process capable of compelling records and sworn testimony can establish whether the failure was merely systemic or something darker. But it is not legitimate to assert intentional enabling without the specific evidentiary markers that would justify that claim.

This is precisely why the “let-it-happen” hypothesis strengthens, rather than weakens, the argument for a coercive inquiry. Where secrecy dominates, where agencies cannot speak freely, and where the public sees multiple warning indicators coupled with catastrophe, suspicion becomes inevitable. The remedy is not to scold suspicion out of existence. The remedy is to build a process that can adjudicate it—by compelling documents, mapping decisions, testing who knew what and when, examining whether warnings were received and acted upon, and determining whether non-intervention was the product of lawful triage, institutional dysfunction, or something worse. Without that compulsion, the truth—whatever it is—remains inaccessible, and the vacuum will continue to be filled by the only thing the public can produce when it lacks evidence: inference, conjecture, and the persistent belief that the most powerful institutions in society are never merely incompetent, but always purposeful.

Chapter 13 — Why Governments Resist Royal Commissions

When governments refuse a Royal Commission after a national trauma, they rarely say, “We don’t want scrutiny.” They speak instead in the language of **efficiency**, **unity**, and **process**. In the Bondi aftermath, the Commonwealth’s refusal to establish a federal Royal Commission was defended publicly on familiar grounds: that Royal Commissions take years, that the moment demands urgency rather than protracted inquiry, and that public hearings risk inflaming division or exposing sensitive national security material that does not lend itself to open examination. These reasons are not invented. They are often real considerations in democratic governance. Yet they are also incomplete, because they omit the strategic political calculus that almost always sits behind the public script. To understand why governments

resist Royal Commissions, one must map both layers: what is said, and what is plausibly driving the refusal even when it remains unsaid.

The first public reason is **speed**. The government has argued that a Royal Commission would take years to deliver findings and that the public needs answers quickly. In this framing, a review process is treated as the rational compromise: it can access classified material, report to government, and produce recommendations on a faster timeline. Speed matters, but it also provides a convenient shield, because it permits the government to pose refusal as responsible management rather than as resistance to scrutiny. A Royal Commission is painted as slow and theatrical; a review is painted as agile and practical. The difficulty is that speed does not answer the core legitimacy problem raised by Bondi: not merely “what reforms should occur,” but “what happened, who knew, what was requested, what was decided, and why.” Where the public suspects a broken chain of accountability, a faster answer is not necessarily a more trusted answer.

The second public reason is **social cohesion**. The Prime Minister has explicitly framed the decision as choosing “unity and urgency rather than division and delay.” This is politically effective language because it casts critics of the refusal as preferring conflict over healing. But cohesion arguments also perform another function: they shift the debate away from institutional accountability and toward community management. In a case involving antisemitic violence, a Royal Commission can be characterised as risking the amplification of extremist rhetoric, the validation of hateful narratives, or the provocation of further community tension. The Attorney-General has advanced precisely this concern, warning that a Royal Commission would “platform” those responsible for “the worst statements” and “revive” those voices publicly. The cohesion rationale may be sincere, but it also narrows what is politically thinkable: it implies that full transparency is itself a threat to stability.

The third public reason is **secrecy and national security evidence**. The Attorney-General has argued that national security information, by definition, does not lend itself to public inquiry. There is truth in this. Intelligence is not merely “information”; it is often inseparable from sources, methods, international sharing arrangements, and ongoing operations. Yet secrecy arguments can be used in two ways. They can be used defensively, to prevent harm. They can also be used tactically, to control the accountability surface. The public cannot assess which function is operating when secrecy is invoked, because secrecy prevents verification. In practice, the secrecy rationale can become a blanket that covers everything from legitimate source protection to simple reputational embarrassment.

Those are the public reasons. Alongside them sit plausible political reasons—rarely admitted, yet entirely consistent with how power behaves in democratic states after crisis.

The first plausible reason is **damage control**. Royal Commissions are unpredictable. They have coercive powers, they generate documents, they compel sworn testimony, and they often produce rolling revelations that dominate news cycles for months or years. A government managing a crisis will prefer controlled processes because uncontrolled processes can detonate not only agency reputations but ministerial credibility. A review can be structured, scoped, and timed; a Royal Commission can expand and expose. Even where government intends genuine reform, it may still resist a Royal Commission because it cannot tolerate the political volatility of unbounded public examination.

The second plausible reason is **precedent management**. Once a government grants a Royal Commission in response to one major failure, it becomes far harder to refuse demands for a Royal Commission in future crises. The refusal is therefore not only about Bondi; it is about protecting the government from becoming permanently subject to the coercive inquiry model whenever a catastrophe occurs. This is not a noble motive, but it is a common one: governments fear establishing a benchmark that reduces their ability to control subsequent crises.

The third plausible reason is **narrative management**. In the immediate aftermath of an atrocity, the government wants to project competence, compassion, and decisiveness. A Royal Commission risks embedding an alternative narrative: one of state failure, ignored warnings, bureaucratic negligence, or institutional disarray. A review process allows the government to emphasise “learning lessons” rather than “assigning blame,” which is often politically safer even when it is institutionally insufficient. Narrative management also intersects with how motive is framed. The more the event is described through the prism of “ISIS-inspired terrorism,” the more the solution space shifts toward intelligence frameworks and powers—territory where governments can claim they are acting. The more it is described through the prism of “antisemitic targeting and inadequate protection,” the more it forces questions about concrete policing decisions and pre-event warnings—territory where governments can be pinned to specific acts and omissions.

The fourth plausible reason is the protection of **institutional relationships**. Commonwealth government depends on ongoing cooperation with security agencies and state counterparts. A Royal Commission can strain those relationships by turning agencies into adversaries under cross-examination, exposing internal dysfunction, and generating a fear culture in which officials become more defensive rather than more candid. Governments often prefer to preserve cooperative relationships and pursue reform through quieter mechanisms rather than risk open conflict with the very institutions they must rely on to keep the public safe.

The fifth plausible reason is **risk displacement**. Royal Commissions, by design, surface uncomfortable trade-offs: resource scarcity, missed escalation, flawed triage, and failures that may implicate multiple agencies across multiple jurisdictions. When these trade-offs are made visible, the public may demand reforms that are politically difficult—massive resourcing, intrusive powers, or structural rearrangements that carry civil liberties implications. A government may prefer a review that can recommend incremental changes without forcing a national confrontation with the costs, risks, and moral compromises embedded in modern security governance.

Seen in this full light, resistance to Royal Commissions is not a single reason but a composite calculus. Governments publicly invoke speed, cohesion, and secrecy because those are defensible rationales that sound responsible. But beneath those rationales sits an equally powerful set of incentives: to contain political damage, avoid precedent, manage narrative, preserve institutional alliances, and keep the accountability surface within predictable limits. The public knows this intuitively, which is why refusal itself often becomes an accelerant of distrust. When citizens suspect that the state is choosing a process that maximises control rather than truth, trust collapses not because people are unreasonable, but because they recognise the logic of power: when power is under threat, it rarely chooses the process that exposes it most.



Chapter 14 — Why Reviews Often Fail the Trust Test

After a national trauma, governments reach instinctively for **reviews**. Reviews are administratively convenient, politically manageable, and rhetorically persuasive. They signal action, promise lessons, and offer a timeline. In the Bondi aftermath, the Commonwealth chose an independent review model rather than a federal Royal Commission, positioning that choice as the responsible path: faster, less divisive, and more compatible with national security secrecy. Yet reviews often fail the public trust test precisely because they sit in the wrong relationship to power. They can generate recommendations without producing legitimacy, because legitimacy is not generated by conclusions; it is generated by **visible accountability**.

The trust test has a simple logic. When the public suspects a broken chain—warnings issued, requests made, a protective posture that appears insufficient, and mass casualties as the

outcome—it does not merely want reform proposals. It wants to see the chain of responsibility traced in daylight: who knew what, when, what decisions were taken, what options were available, and why particular choices were made. A review, particularly a closed one, can report on what it believes happened without enabling the public to verify that belief. That is the first reason reviews fail: they ask citizens to accept a narrative **on trust** in the very moment when trust has collapsed.

A closed review is structurally limited by its method. Its fact-finding is usually conducted in private, through briefings, document access mediated by agencies, and interviews without public cross-examination. That method may be necessary where classified material exists, but it also creates an asymmetry that is fatal to legitimacy: the review can see evidence the public cannot, and then the public is asked to accept the review's summary without access to the evidentiary substrate. In a high-stakes failure, the public does not merely doubt the conclusions; it doubts the process by which those conclusions were reached, because it cannot test the process. Secrecy may be unavoidable in certain respects, but secrecy also functions as an accelerant of suspicion when the subject under examination is the performance of secretive agencies themselves.

The second limitation is that many reviews lack the **coercive architecture** needed to break institutional defensiveness. Even where a reviewer is credible and independent, the inputs still depend on an ecosystem of agencies whose reputations and legal exposures are at stake. In that environment, information-sharing becomes strategic. Agencies tend to provide what is asked for, not necessarily what is most damaging. They interpret questions narrowly, frame decisions as rational trade-offs, and emphasise constraints rather than errors. This does not require deliberate lying; it is the ordinary behaviour of institutions under scrutiny. A Royal Commission, by contrast, is designed to counteract this behaviour through compulsion: subpoena powers, sworn evidence, penalties for non-compliance, and the capacity to force disclosure beyond what agencies would volunteer. Reviews that lack these features can produce accurate conclusions, but they cannot reliably defeat strategic withholding or selective framing. The public senses this intuitively, which is why “independent review” often feels like “internal audit” under another name.

The third limitation is the absence of public testing through **cross-examination** and adversarial challenge. Trust is not restored by an authority figure stating what happened; it is restored when the public sees claims tested, contradictions surfaced, and excuses confronted. Where a community risk assessment is reported to exist, where some officials say they were unaware of it, while political leaders say contact occurred, the dispute is not settled by a written conclusion. It is settled by a process that can compel the relevant documents, compel the relevant witnesses, and test competing accounts in a forum where evasions carry consequences. Reviews rarely provide that forum. They can describe “issues with communication” without forcing clarity on whether a warning was received, who read it, and what was done with it.

The fourth limitation is that reviews are often scoped to avoid “blame,” and that avoidance can itself destroy legitimacy. Governments routinely insist that a review is about learning lessons, not finding fault. The public hears that as avoidance of accountability. A society can accept that complex systems fail. What it struggles to accept is the refusal to name responsibility where responsibility exists. Royal Commissions are not perfect, but they are structurally capable of making explicit findings—about systemic failures, policy errors, and, where warranted, individual decision-making failures. Reviews often soften language,

distribute responsibility into abstractions, and end with recommendations that everyone can endorse without anyone being directly accountable. That pattern satisfies administrative governance but fails democratic legitimacy.

The fifth limitation is the credibility problem created by who controls timing and release. Reviews are commissioned by government, designed by government, and reported to government. Even where the reviewer is respected, the public cannot escape the perception that government controls the boundaries. A Royal Commission, though established by government, has a more autonomous posture once created. It has public hearings, independent counsel assisting, and a procedural architecture that makes “massage” more difficult. By contrast, reviews can be delayed, partially released, summarised, or framed through ministerial statements. That is why reviews struggle to counter the impression of narrative management: the very fact of government commissioning becomes part of the suspicion.

The sixth limitation is that reviews often operate in parallel with criminal investigations and inquests, which can be used—legitimately or tactically—to justify silence. After Bondi, NSW Police have declined detailed comment in reporting, referring to ongoing critical incident investigation, criminal investigation, and coronial processes. Those processes matter, but they also create a common pattern: key operational decisions are never publicly ventilated while they are most relevant, and by the time they can be ventilated, political urgency has faded. A review produced during that period can only work with partial visibility or agency-provided accounts, and the public knows it. The consequence is that the review becomes another layer of institutional communication rather than a truth-forcing mechanism.

The result is a predictable legitimacy gap. Reviews can produce reform lists and policy adjustments. They can recommend better liaison, improved information sharing, refined thresholds, and enhanced training. They can do all of this and still fail to restore trust, because the public’s crisis is not purely technical. It is epistemic: the public no longer believes it has access to the truth about what the state did and did not do. Where epistemic trust collapses, the remedy is not only competence; it is **transparency under compulsion**.

This is why the “trust test” matters. In a case like Bondi, the central question is not whether government can produce an answer quickly; it is whether government can produce an answer in a way that the public will accept as legitimate. Closed reviews frequently fail because they are structurally incapable of demonstrating the chain of accountability to the public. They may generate conclusions. They may even generate correct conclusions. But without public testing, coercive access to records, and a process that forces institutions to disclose what they would prefer to keep sealed, conclusions become another form of assertion. And when trust has collapsed, assertion is the one thing that no longer works.



Chapter 15 — The Case for a Royal Commission to Restore Public Trust

A Royal Commission is not a ritual of outrage. It is a constitutional instrument designed for a specific democratic purpose: to restore legitimacy when the public can no longer determine whether the state has told the truth about a catastrophe. In the Bondi case, the demand for a federal Royal Commission did not arise merely because the event was horrific. It arose because the public record contains multiple reported warning indicators, overlapping jurisdictions, and competing claims about who knew what and when. When such conditions exist, trust collapses not because citizens are irrational, but because they recognise a structural reality: institutions under scrutiny cannot credibly investigate themselves without coercive oversight, and closed processes cannot repair legitimacy in a crisis defined by informational asymmetry.

The democratic rationale begins with a simple proposition. Modern security governance is built on secrecy. Intelligence agencies, counterterror operations, and even parts of event policing operate behind layers of confidentiality. In ordinary times, democratic societies accept this, trading transparency for protection. But when an attack occurs in circumstances where warnings were allegedly communicated and protection was allegedly sought, secrecy becomes politically combustible. It converts uncertainty into suspicion. It creates a vacuum in which citizens cannot evaluate whether failure occurred, whether it was preventable, and whether official explanations are complete. A Royal Commission is the mechanism by which a democracy temporarily reverses that normal relationship: it compels the state to show its working to an independent authority, and, so far as lawful and safe, to the public itself.

Bondi has generated precisely the kind of contested chain that a Royal Commission is built to examine. Reporting has described a pre-event risk assessment by a community security organisation warning of high risk, including violent antisemitism, weeks before the attack. Reporting has described requests for increased police protection and political acknowledgment that police presence was “clearly not enough.” Reporting has also described

that one alleged perpetrator had previously been known to ASIO in connection with extremist networks years earlier, before that attention was closed. Each element, viewed alone, can be explained away: warnings may not have been received; resources may have been constrained; intelligence interest may have been historic and non-actionable. But in combination, these elements create a legitimacy crisis, because the public can plausibly believe that the system had enough signals to treat the event as high risk. When the state cannot demonstrate the chain of reasoning that led to its protective posture, suspicion becomes the default.

The Royal Commission case is therefore fundamentally about epistemic authority: who has the power to determine what happened. In a democracy, the state cannot be the sole arbiter of its own performance after catastrophe, because the state has direct conflicts of interest. It has reputational interests, legal exposure, and political incentives to frame failure as unforeseeable. Even well-meaning agencies tend to interpret events through self-protective lenses. This is not an allegation of conspiracy; it is institutional psychology. The Royal Commission exists because democracies recognise that, without an external truth-forcing mechanism, the same institutions that failed can also control the narrative of why they did not fail. Trust cannot be restored on that basis.

What makes a Royal Commission unique is coercion paired with independence. It can compel documents that would not otherwise be produced. It can compel sworn testimony and test it. It can reconcile contradictions between agencies, between government and community organisations, and between official accounts and documentary timelines. It can establish a chronology that is not built from press statements but from verifiable records. It can identify where information moved and where it died, not by inference but by tracing actual communications, logs, briefings, and decision points. It can also expose failures that are otherwise invisible precisely because they are omissions: warnings not escalated, requests not logged, briefings not sent. These are the failure modes that reviews often cannot surface, because omission leaves little trace unless someone is compelled to explain the absence.

The democratic value of coercive fact-finding is not only retrospective truth. It is prospective legitimacy. When the public believes the state has concealed or sanitised the causes of failure, it does not merely distrust the specific institutions involved; it begins to distrust the entire governance system—policing, intelligence, courts, and political leadership. That distrust is itself a national security problem, because it corrodes cooperation, fuels radicalisation narratives, and fractures social cohesion. A Royal Commission therefore functions as a stabilising instrument: it offers a pathway for grief and anger to be converted into public truth rather than into permanent cynicism.

The Commonwealth's argument against a Royal Commission emphasises speed, unity, and secrecy. Those considerations matter, but they do not resolve the core legitimacy question raised by Bondi. A fast answer that the public does not trust is not a solution; it is another accelerant. Unity achieved by suppressing scrutiny is not unity; it is merely quiet distrust. Secrecy justified by national security may be necessary in parts, but the broader decision chain—receipt of warnings, resourcing decisions, liaison protocols, and governance accountability—can be examined without revealing operational tradecraft, especially if the inquiry is properly structured to protect sensitive sources. The point is not to publish everything. The point is to ensure that an independent body can see everything and can provide the public with findings that are demonstrably grounded in compelled evidence rather than curated narratives.

A Royal Commission is therefore justified when three conditions converge. First, there are multiple warning indicators that, if true, raise the possibility that harm was preventable. Second, responsibility is distributed across multiple agencies and jurisdictions, creating a risk of buck-passing and narrative divergence. Third, key operational facts cannot be resolved through normal public processes because of secrecy claims or ongoing investigations. Bondi appears to satisfy each condition, which is why families and community voices insist that a review is not enough. They are not asking for perfection from the state. They are asking for the one thing that makes democratic life sustainable after catastrophe: an independent, coercive process that can establish the truth of what happened and make accountability real rather than rhetorical.

The final democratic rationale is the simplest. When a government refuses the strongest truth-finding mechanism available, it must accept the political consequence: the refusal itself becomes evidence, in the public mind, that the state fears what would be uncovered. That inference may be unfair. It may be wrong. But it is predictable, because citizens understand how power behaves when threatened. A Royal Commission, properly designed, is the mechanism by which a democracy prevents that inference from calcifying into permanent distrust. It restores public trust not by demanding belief, but by earning it through forced disclosure, sworn evidence, and a record that can withstand the harshest scrutiny.

Chapter 16 — Terms of Reference That Actually Answer the Hard Questions

A Royal Commission can be made useless by being made vague. It can become a theatre of generalities—“lessons will be learned”—without forcing clarity on the only questions that matter after a preventable disaster: **who knew what, when did they know it, what was requested, what was decided, what was resourced, and why.** The Bondi case requires Terms of Reference that are tight enough to compel an auditable chronology, yet broad enough to capture the full end-to-end chain across Commonwealth and NSW systems. The objective is not to criminalise hindsight or to demand omniscience. The objective is to build a factual record that allows the public to understand whether blind spots were unavoidable or whether they were the product of structural failures, poor judgments, or governance gaps.

The scope must be designed around five principles. First, the inquiry must be empowered to obtain and examine **all relevant records**, including classified records, subject to appropriate protective procedures. Second, it must reconstruct a single integrated chronology, not separate agency narratives. Third, it must identify decision points and the rationales for decisions, not merely describe outcomes. Fourth, it must distinguish between what was known, what was inferred, and what was assumed. Fifth, it must produce findings capable of being understood publicly without compromising legitimate national security sensitivities.

What follows is a Terms of Reference structure drafted to force those outcomes.

1. Chronology and Knowledge Mapping

- (a) The Commission shall establish a comprehensive chronological account from 1 January 2019 to 31 December 2025 of all material information held by, provided to, or accessible to relevant agencies concerning Naveed Akram, the second alleged offender, any known associates, and any relevant persons or entities connected to the

planning, motivation, financing, procurement, or facilitation of the Bondi terror attack.

(b) The Commission shall identify, for each material item of information, the date of acquisition, the agency or unit holding it, the classification applied, the assessed reliability, and whether and when it was disseminated to any other agency or decision-maker.

2. ASIO Engagement, Assessment, and Closure Decisions

(a) The Commission shall inquire into the basis, scope, and outcomes of ASIO's prior engagement with Naveed Akram in 2019, including the criteria applied in assessing threat, the reasons for any decision to deprioritise or close attention, and any subsequent review or re-scoring of risk.

(b) The Commission shall determine whether any later indicators or intelligence holdings should reasonably have triggered reassessment, escalation, or renewed attention, and if not, why not.

(c) The Commission shall examine the adequacy of ASIO's triage frameworks and closure governance, including documentation standards, supervisory oversight, and audit mechanisms.

3. Community Risk Assessments and Warning Pathways

(a) The Commission shall inquire into the creation, content, dissemination, and receipt of any community risk assessments, including any assessment dated 26 November 2025, and any other warnings relating to heightened antisemitic risk, lone-actor risk, or terrorist threat affecting the Chanukah gathering at Bondi.

(b) The Commission shall determine precisely who within NSW Police Force, NSW Government, and any Commonwealth agency received such material, when it was received, how it was recorded, and what assessment was made of it.

(c) The Commission shall assess whether existing warning pathways were fit for purpose, including whether formal acknowledgement and escalation protocols existed and were followed.

4. Requests for Police Protection and Operational Decision-Making

(a) The Commission shall inquire into any requests made by organisers, community representatives, or security bodies for enhanced police presence or protective measures for the Bondi Chanukah event, including the timing, content, recipient, and documentation of those requests.

(b) The Commission shall determine what operational plan existed for the event, including staffing levels, placement, command structure, escalation arrangements, and contingency planning, and shall determine what changes, if any, were considered or rejected in response to risk assessments or requests.

(c) The Commission shall identify the reasons for the resourcing and deployment decisions made, including the criteria used to rate risk, competing operational demands, and whether any decision-maker concluded that resourcing was insufficient or accepted an elevated risk.

5. Inter-Agency Information Sharing and Silo Effects

(a) The Commission shall examine the information-sharing arrangements between ASIO, AFP, NSW Police Force, NSW Government agencies, and firearms licensing authorities relevant to terrorism risk, event protection, and firearms access.

(b) The Commission shall identify points at which information failed to move, was delayed, was compartmentalised, was downgraded, or was not acted upon, and shall determine whether such failures were caused by legal constraints, policy settings, culture, capability limitations, or governance weakness.

(c) The Commission shall assess whether structural incentives existed to withhold,

minimise, or narrowly interpret information, including reputational incentives, liability concerns, and the avoidance of discoverable records.

6. Firearms Licensing and Household Risk Management

(a) The Commission shall inquire into the firearms licensing status of any relevant household, the storage compliance regime, inspection practices, and the adequacy of controls designed to prevent access by unlicensed persons.

(b) The Commission shall examine whether terrorism-related assessments, investigations, or warnings were capable of being considered within licensing and suitability determinations, including household proximity issues, and if not, why not.

(c) The Commission shall assess whether any changes implemented after the attack indicate that prior settings were insufficient to manage foreseeable household access pathways.

7. Travel, Networks, and Investigative Triggers

(a) The Commission shall examine the relevance of any overseas travel by alleged offenders in the period 1 October 2025 to 14 December 2025, including what Australian agencies knew, what triggers were engaged (if any), and what action was taken.

(b) The Commission shall determine whether any contact patterns, associations, or ideological influences were known or reasonably discoverable, and whether the system's thresholds for escalation were fit for purpose.

8. Motive Framing and Public Communication

(a) The Commission shall examine how motive was characterised in official communications, including the relative emphasis on antisemitic targeting versus ISIS-inspired ideology, and whether any framing choices affected operational accountability, public understanding, or reform priorities.

(b) The Commission shall assess whether public communications were accurate, sufficiently transparent, and appropriately balanced against security constraints.

9. Preventability, Opportunities, and Counterfactual Analysis

(a) The Commission shall determine, on the balance of probabilities, whether there were reasonable opportunities for intervention, disruption, or enhanced protection that could have reduced harm, and shall identify the most significant missed opportunities, if any.

(b) The Commission shall distinguish between failures of law, failures of policy, failures of judgment, failures of resourcing, and failures of execution.

10. Recommendations and Accountability Mechanisms

(a) The Commission shall make recommendations to strengthen end-to-end governance across warning receipt, escalation, event protection triage, intelligence closure discipline, and firearms household risk management.

(b) The Commission shall recommend permanent accountability mechanisms, including mandatory documentation standards, independent audit powers, and clear escalation protocols for high-risk community events.

The virtue of these Terms of Reference is that they do not allow escape into abstraction. They force decision-makers to answer concrete questions, supported by documents and sworn evidence. They also force the inquiry to produce a record capable of public comprehension: not merely that “gaps existed,” but exactly where they existed and why. That is what restores public trust—clarity that cannot be produced by reassurance, and legitimacy that cannot be produced by closed conclusions.

Chapter 17 — From Failure to Reform: What Changes Without Breaking the Rule of Law

The political temptation after terror is always the same: expand powers, widen surveillance, harden the state. That reflex is understandable, but it carries its own long-term danger. If every failure is met with permanent emergency settings, the society may reduce one category of risk while eroding the conditions that make democratic life worth defending. The disciplined response to Bondi, therefore, is not the maximalist response. It is reform that reduces **blind spots** and increases **accountability** without normalising indiscriminate monitoring, indefinite suspicion, or broad coercive powers untethered from clear thresholds. The central challenge is to build a system that is harder to surprise, without becoming a system that treats the entire public as suspect.

The first reform priority is not more information; it is better **information movement**. Bondi's defining risk was fragmentation—warnings, triage decisions, and risk assessments distributed across separate nodes. The remedy is a formal, auditable chain for high-risk event warnings. Where a community security organisation provides a written assessment rating an event as high risk, the system should require a mandatory acknowledgement, a logged escalation, and a documented decision within a set period. This is not a new surveillance power. It is governance discipline. It forces agencies to treat warnings as actionable objects rather than as informal background noise, and it creates an evidentiary trail that can later be audited without disclosing sensitive tradecraft. It also protects police and government from unfair insinuation, because they can prove what they received and how they responded.

The second priority is to redesign event policing triage around **targeted-community risk** as a first-order variable. Event policing often treats faith gatherings, cultural festivals, and community events as low-complexity unless there is a specific threat. The Bondi experience shows why this model is insufficient in a heightened antisemitism environment. The reform is not to police every event heavily. The reform is to mandate a risk rubric that weights targeted-group threat levels and symbolic dates as triggers for uplift decisions, and requires decision-makers to document why an uplift was granted or denied. This reform strengthens protection while preserving the rule of law, because it operates through transparent criteria and accountable decision records rather than secret expansions of power.

The third priority is to harden the **closure governance** of intelligence work. No agency can monitor everyone indefinitely, but closure must be treated as a risk decision that carries ongoing accountability. When a subject has previously intersected with terrorism-related investigation, closure should be accompanied by a structured "revisit logic": clearly defined triggers that would require reassessment, and periodic review where those triggers exist. This does not require perpetual surveillance. It requires disciplined documentation, supervisory oversight, and independent auditing of closure decisions—particularly where later harm would predictably generate public shock. Done properly, it reduces the closure problem without creating a permanent suspect class, because it operates through narrow, defined criteria and internal governance rather than mass collection.

The fourth priority is to align firearms licensing with **household risk** while preserving fairness. Bondi has already prompted NSW reforms that recognise a basic truth: firearms risk is not purely individual; it is environmental. The rule-of-law version of this reform is not

“deny licences by association” in an arbitrary way. It is to build a narrowly defined, reviewable regime where terrorism-related investigation status—whether of an applicant or a household member—can trigger enhanced scrutiny, storage verification, and, in defined circumstances, refusal or suspension with procedural fairness protections. The objective is not punishment without conviction; it is risk mitigation through proportionate safeguards that are contestable and reviewable, rather than discretionary and opaque.

The fifth priority is to create a **single operational risk picture** for high-risk events that integrates community assessments, policing posture, and relevant intelligence constraints without collapsing everything into one super-database. The rule-of-law way to do this is not indiscriminate data pooling. It is a joint risk conference mechanism for nominated high-risk events—time-limited, event-specific, and auditable—where the relevant parties can share what is necessary to decide resourcing and posture. The mechanism should emphasise minimisation: share only what is needed, retain only what is necessary, and ensure an independent audit function can verify that disclosure and retention were proportionate.

The sixth priority is to strengthen accountability without weaponising it. Agencies often become defensive after attacks because accountability processes are experienced as blame factories. The solution is to separate **operational learning** from **disciplinary consequences** while still enabling both where warranted. A permanent independent audit office—empowered to inspect warning receipt logs, triage decisions, and information-sharing pathways—can identify systemic gaps without requiring a Royal Commission every time. Yet, where negligence or misconduct is identified, pathways for consequences must exist. This preserves the rule of law by ensuring accountability is not arbitrary and by protecting individuals from scapegoating while still making the system answerable.

The seventh priority is to fix the public communication problem, because legitimacy is partly built in language. Framing choices after attacks—whether emphasising antisemitic targeting or ISIS inspiration—shape what reforms feel justified. The reform here is a standard of candour: early communications should distinguish clearly between what is known, what is suspected, and what is still being assessed, and should avoid rhetorical shortcuts that flatten accountability into slogans. This is not a cosmetic reform. It reduces speculation by reducing the informational vacuum that speculation feeds on.

The eighth priority is to design reforms with **sunset and review** logic so that emergency responses do not become permanent. Where new powers are proposed, they should be narrow, time-limited, and subject to independent review with public reporting. Where surveillance settings are expanded, they should be matched with strong minimisation, warrant thresholds, and enforceable oversight. A democracy does not preserve itself by refusing all new powers; it preserves itself by refusing to make “temporary fear” the justification for permanent state expansion.

The final point is that reform must be measured against the correct benchmark. The benchmark is not whether the state can promise “this will never happen again.” No honest system can promise that. The benchmark is whether the state can demonstrate that it has reduced preventable blind spots without dissolving liberty into permanent emergency governance. Bondi demands a system that is harder to surprise and easier to audit; more disciplined in how warnings are received and acted upon; more coherent in how intelligence closure decisions are governed; more realistic about household firearms access risk; and more transparent about how protection decisions are made. None of that requires a surveillance

society. It requires what democratic security has always required, but too often fails to deliver: **structured accountability**, **proportionate prevention**, and a rule-of-law state willing to expose its own decision chains to independent scrutiny when catastrophe suggests those chains may have broken.

Chapter 18 — The Watchlist Paradox: When Prior Interest Becomes Post-Attack Evidence

The most corrosive feature of modern counterterrorism is not that agencies miss things. It is that, after an attack, the public discovers that the alleged perpetrator had previously been known to the system—investigated, assessed, and then deprioritised or closed. That sequence creates a paradox. Before an attack, “no ongoing threat” is treated as a rational conclusion under uncertainty. After an attack, the same conclusion becomes retroactive evidence that the system is either incompetent, evasive, or worse. Bondi fits this pattern in a way that makes the legitimacy shock almost inevitable. Albanese has confirmed that Naveed Akram first came to ASIO’s attention in 2019, was investigated for several months, and was later assessed as posing no ongoing threat. After the attack, Albanese also acknowledged that something was missed in the period between that earlier closure and the eventual violence. In a society already sensitised to rising antisemitism and warnings around public Jewish events, that is the worst possible combination: prior interest exists, closure is asserted, and yet the attack still happens.

The public’s response is not irrational. It is based on a commonsense inference about risk: if an individual has ever crossed the threshold of intelligence attention for extremist association, then certain later factors should be treated as amplifiers. Overseas travel shortly before an event, contact patterns that suggest ideological reinforcement, and household access pathways to lethal capability are not proof of intent, but they are precisely the kinds of indicators that should trigger heightened scrutiny and a hardened protective posture. Bondi’s case has been widely discussed in those terms, including the focus on November travel to the Philippines and the household firearms-access question that follows whenever the alleged offender is not the licensed owner but has proximity to a licensed environment. The point is not that any one factor is determinative. The point is that when factors accumulate, the public assumes the system should respond differently.

This is the watchlist paradox. Intelligence agencies defend closure decisions by pointing to legal thresholds, volume pressures, and the reality that countless individuals are assessed and never offend. That is true. But it is also true that, in high-impact attacks, the public does not accept system arithmetic as a sufficient explanation. The public cares about low-frequency, high-consequence risk, and it expects the system to behave differently when the potential outcome is mass casualty. This mismatch in risk philosophy is one of the deepest drivers of distrust. Agencies triage for probability; citizens triage for catastrophe.

International precedent shows this pattern is not unique. Multiple democracies have recorded cases in which perpetrators were previously investigated and then later committed attacks after cases were closed or downgraded. The recurring lesson from those precedents is not that agencies intentionally “let it happen.” The lesson is that closure decisions are often governed as if they were administrative endpoints rather than risk decisions with continuing accountability. Once a case is closed, it tends to fall out of the decision ecology unless a clear

trigger forces reassessment. Yet in real life, radicalisation can be nonlinear and opportunistic. People can sit in low-visibility ideological space for years and then accelerate rapidly in response to events, networks, or personal crises. A system built around closure as an endpoint will always be vulnerable to that acceleration.

The Bondi question, therefore, is not whether prior ASIO attention “proves” preventability. It is whether the system’s closure and re-opening logic is fit for purpose in the modern threat environment. If a subject has previously been of terrorism interest, what is the governance discipline that ensures reassessment occurs when risk amplifiers appear? If travel to sensitive regions occurs shortly before a high-risk symbolic event, what is meant to happen in the system—who is notified, what thresholds apply, and what protective posture is triggered? If a household contains lawful firearms and also contains a person previously assessed for extremist links, what safeguards exist to prevent foreseeable access pathways? These are not rhetorical questions. They are governance questions, and they cannot be answered by press conferences.

This is also why the watchlist paradox becomes politically explosive. Governments tend to treat these matters as national security questions best handled in closed reviews, because the underlying material is classified and the judgments are contestable. But closed reviews rarely restore trust, because they ask the public to accept a conclusion without seeing the decision chain that produced it. When the public already believes the system had prior interest and still failed, “trust us” is not a remedy; it is an accelerant. The only process that can neutralise the watchlist paradox is one that compels records, compels sworn testimony, and reconstructs the full chronology across agencies—what was known, what was shared, what was requested, what was decided, what was resourced, and why.

The argument for a Royal Commission, in this context, does not depend on proving Machiavellian intent. It depends on a narrower and stronger democratic claim: where an alleged perpetrator was previously known to intelligence, and where subsequent risk amplifiers appear to have existed, the public is entitled to a coercive, independent account of how the system reasoned and where it failed. Without that account, Bondi becomes another entry in the international pattern: prior interest, closure, catastrophe, and then a permanent residue of distrust—because the state asks citizens to believe in competence while refusing to show the evidence that competence would produce.

Chapter 19 — The Jurisdictional Shield: How the “NSW Royal Commission + Federal Review” Route Dilutes Commonwealth Scrutiny

When a terror event triggers demands for a Royal Commission, the decisive issue is not rhetorical commitment to “accountability,” but the structure of the process that will be allowed to examine the state. Bondi has exposed a familiar manoeuvre in federations: channel the most coercive inquiry mechanism into the **state** lane, while keeping **Commonwealth** intelligence and federal law enforcement inside a **review** lane. The effect is not the total absence of scrutiny of ASIO and the AFP, but a form of scrutiny that is structurally less public, less coercive, and less capable of producing the kind of legitimacy reset that only open, adversarial fact-finding can deliver.

The jurisdictional reality is straightforward. A NSW Royal Commission is inherently strongest when examining NSW institutions: NSW Police Force operational planning, resourcing decisions, ministerial briefings within NSW, and the handling of community warnings inside NSW government channels. It is weaker when the chain of responsibility runs into Commonwealth systems because Commonwealth agencies sit under different legal authority, different secrecy settings, and different accountability pathways. A state Royal Commission can request cooperation and can draw adverse inferences from non-cooperation, but it is not the same as an inquiry established with Commonwealth jurisdiction—or a genuinely joint commission—where Commonwealth agencies are unambiguously inside the coercive perimeter from the start.

This is why the “NSW Royal Commission + federal review” route matters. In public terms, it can be presented as the best of both worlds: NSW runs a Royal Commission, the Commonwealth runs an independent review, and everyone cooperates. In legitimacy terms, it is not the best of both worlds. It is a partitioning strategy. It separates the domain where public anger is most acute—policing posture, warnings, and protection of a Jewish event—into a coercive state process, while placing the most sensitive and reputation-critical questions—intelligence triage, closure decisions, watchlist governance, and federal counterterrorism thresholds—into a process that is typically more controlled and less publicly ventilated.

This partition has three practical consequences. First, it reduces **public visibility** of Commonwealth agency decision-making. Reviews, even independent ones, tend to operate in closed settings because their evidentiary substrate is classified. They may access material that the public will never see. They may reach conclusions that cannot be fully explained. The public is then asked to accept a summary, and in a trust deficit environment, summaries do not restore legitimacy. By contrast, Royal Commissions—although they can hold closed sessions for sensitive evidence—are built to do their work in public wherever possible. The difference is not cosmetic. It is the difference between assertion and demonstration.

Second, the partition reduces **adversarial testing** of Commonwealth narratives. A Royal Commission has the capacity to compel sworn evidence and test it through counsel assisting and, where permitted, through questioning that exposes contradiction and evasion. A review is typically not structured as a quasi-judicial forum. Even if the reviewer is respected, the method is often consultative and document-based rather than adversarial. This changes the incentives. In adversarial settings, institutions are forced into precision. In consultative settings, institutions can frame decisions as reasonable trade-offs, and the reviewer may have limited procedural capacity to force disclosure beyond what is offered or to test credibility through public contradiction.

Third, the partition enables **accountability displacement**. When a failure spans multiple agencies, each institution can point to the other: police can imply they lacked intelligence; intelligence can imply there was no actionable threshold; licensing can imply it is not an intelligence function; ministers can imply agencies did not brief them; agencies can imply ministers did not ask. A single coercive inquiry with cross-jurisdiction reach is designed to defeat this by compelling the integrated chain. A split model—NSW RC for NSW actors, review for Commonwealth actors—risks preserving the very silo structure that created the failure in the first place. The public ends up with two narratives, two chronologies, two sets of findings, and no single forum that compels the complete integrated truth.

This is why, even if Albanese's stated rationale is speed and unity, the practical effect is to lessen the intensity of scrutiny applied to ASIO and the AFP compared with what would occur under a federal or joint Royal Commission. It is not that ASIO and the AFP escape scrutiny entirely; it is that their scrutiny occurs in a format that is less public, less coercive, and less legitimacy-restoring. In politics, effects often matter more than stated motives. If the effect is reduced exposure of Commonwealth agencies to the kind of open, document-compelling, sworn-evidence process that could reveal uncomfortable chains of decision, then the public will naturally infer that institutional protection is part of the calculus—even if it cannot be proven without access to internal deliberations.

The democratic danger of this route is not only that it may leave some questions unanswered. It is that it leaves the public with the sense that the state has chosen the architecture of inquiry to control the narrative rather than to expose the full truth. Once that belief takes hold, it metastasises. Every future assurance from government becomes less persuasive. Every future refusal to disclose “operational details” is interpreted as evasion. The legitimacy crisis deepens, not because citizens demand perfection, but because they recognise the logic of jurisdictional partition: when the most powerful agencies are insulated from the most coercive public process, accountability becomes partial by design.

For Bondi, the core accountability claim is therefore structural. If the alleged perpetrators were on the radar of Commonwealth systems at any stage, if the attack intersected with issues of watchlist governance, closure discipline, federal-state information sharing, and counterterror thresholds, then a process that cannot fully compel Commonwealth agencies in the same forum as state agencies will struggle to answer the hard questions. A NSW Royal Commission can expose NSW failures and will be valuable for that. But it cannot, on its own, deliver the integrated democratic reckoning that the public seeks when a national terror event sits at the intersection of Commonwealth intelligence and state policing. That integrated reckoning is precisely what a federal or joint Royal Commission is designed to provide, and it is precisely what this route, by its structure, risks avoiding.